



A psychomechanical study of rolling sounds

Christophe N.J. Stoelinga

► To cite this version:

Christophe N.J. Stoelinga. A psychomechanical study of rolling sounds. Engineering Sciences [physics]. ENSTA ParisTech, 2007. English. NNT: . pastel-00002368

HAL Id: pastel-00002368

<https://pastel.hal.science/pastel-00002368>

Submitted on 20 Apr 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Psychomechanical Study of Rolling Sounds

Christophe N.J. Stoelinga

The work described in this thesis has been carried out under the auspices of the J. F. Schouten School for User-System Interaction Research.

Cover design: R.-P. Kruizinga

CIP-DATA LIBRARY TECHNISCHE UNIVERSITEIT EINDHOVEN

Stoelinga, Christophe N.J.

A psychomechanical study of rolling sounds / by Christophe Norbert Jos Stoelinga.

- Eindhoven : Technische Universiteit Eindhoven, 2007. – Proefschrift. -

ISBN 978-90-386-0914-0

NUR 927

Keywords: Acoustics / Psychomechanics / Rolling sounds / Psychophysics /
Bouncing balls / Numerical modeling

A Psychomechanical Study of Rolling Sounds

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven, op gezag van de
Rector Magnificus, prof.dr.ir. C.J. van Duijn, voor een
commissie aangewezen door het College voor
Promoties in het openbaar te verdedigen
op dinsdag 3 april 2007 om 16.00 uur

door

Christophe Norbert Jos Stoelinga

geboren te Eindhoven

Dit proefschrift is goedgekeurd door de promotoren:

prof.Dr. A.G. Kohlrausch

en

prof.dr. A. Chaigne

Copromotor:

dr. D.J. Hermes

Acknowledgements

Writing a thesis is a long process in which one can use some support from many sources. The unpredictability of results might be one of the principal reasons to choose a future in research, at some point it seems to be impossible to align these results to tell a logical story that fits in a thesis. It is like inverse spaghetti cooking: when the results are hot, they are a mess but once they are cold, they should be tightly organized in a package.

This process was only successful due to my supervisors, Dik Hermes and Armin Kohlrausch, who always kept overview and helped me sometimes at the cost of their free time. They were not afraid to look a little beyond the scope of their own expertise. Special thanks go to Antoine Chaigne for receiving me at ENSTA. Without the stay at his lab, this work could not have been the interdisciplinary work which it is now. I am grateful to French embassy in The Netherlands who made this visit possible with a grant. Not listed as my formal supervisors, but very helpful on many occasions was Mico Hirschberg.

I was able to share, my own expertise, with some undergraduate students who often, in fact, turned out to be also very helpful to me. Especially Bert Peeters and Koen Steegs, who can find some of their results in this thesis (see references in the footmarks of the chapter titles).

The proof that even random processes can result in most preferable results was given by my roommates, who always turned out to be most pleasurable company. Sequentially they were Maxim Scram, Nicolas Quagebeur and Wouter van den Hoogen.

Last but not least I want to thank my parents, and two sisters, Mariëlle and Josine, for providing me a pleasurable youth and my first experiences with rolling and bouncing balls.

Contents

Acknowledgements	v
1. General Introduction	1
1.1 Motivation	1
1.2 Everyday sounds	3
1.3 Information in sound	5
1.4 Psychomechanics	5
1.5 Acoustics	6
1.6 Measurement	7
1.7 Psychomechanics research scheme	8
1.8 Literature	10
1.9 Goals and main contributions of this thesis	13
1.10 Outline	15
2. Bouncing of a Ball on a Plate: A Model and Measurements Based Upon the Mechanical Driving-Point Impedance of the Plate	17
2.1 Introduction	17
2.2 Theoretical analysis	20
2.3 Measurements	24
2.4 Discussion and conclusion	30
3. The Bouncing of Balls: Influences of Spectral and Temporal Variations on Perception	35
3.1 Introduction	35
3.2 Mechanical analysis	36
3.3 General methods	38
3.4 Experiment 1	42
3.5 Experiment 2	44
3.6 Discussion and conclusion	48

4. Temporal Aspects of Rolling Sounds: A Smooth Ball Approaching the Edge of a Plate	49
4.1 Introduction	49
4.2 Measurement setup	51
4.3 Acoustical analysis	52
4.4 Reflections	54
4.5 Wave velocity	58
4.6 Comb-filter model for reflections at the edge	60
4.7 Synthesis	65
4.8 Repetition pitch of echoed noise	66
4.9 Discussion	67
5. Influence of Wave Reflections at the Plate Edge on Perceiving the Rolling Direction of a Rolling Ball	69
5.1 Introduction	69
5.2 Experiment	71
5.3 Synthesis	71
5.4 Method	73
5.5 Stimuli	73
5.6 Results	75
5.7 Control experiment	75
5.8 Results	76
5.9 Discussion	76
6. Time-domain Modeling and Simulation of Rolling Objects	79
6.1 Introduction	80
6.2 Single impact on a damped plate	82
6.3 Adaptation of the model for moving sources	85
6.4 Simulations	92
6.5 Comparison between measured and simulated plate acceleration	100
6.6 Conclusion	106
6.A Numerical scheme	109
6.B Constant moving force	110
6.C A simple non-linear system	112
7. Using Psychometric Functions to Determine the Auditory Capability to Perceive the Size and Speed of a Rolling Ball	121
7.1 Introduction	122

7.2	Experiment 1: Paired comparisons	126
7.3	Comparison with other studies	132
7.4	Experiment 2: Absolute magnitude estimation	137
7.5	Discussion	145
7.A	Psychological scales	151
8.	Using Multidimensional Scaling to Determine Distances in the Percep- tual Space of Rolling Sounds	155
8.1	Introduction	156
8.2	Stimuli	158
8.3	Method	159
8.4	Results	159
8.5	Discussion	165
8.A	Transformation of the MDS space	168
9.	Discussion	171
9.1	Introduction	171
9.2	Psychomechanic research scheme	174
	Bibliography	183
	Summary	191
	Curriculum Vitae	195

1 | General Introduction

1.1 Motivation

The research reported in this thesis is concerned with the mechanical and the perceptual processes that take place when we listen to the sounds produced by objects bouncing and rolling over a plate. Bouncing and rolling sounds are chosen because, by listening to these sounds, the listeners can form clear estimates of the physical properties of the moving objects such as their size and their speed, or the material from which they are made. This makes it possible to study the relation between the mechanical parameters that determine the acoustic properties of the sound radiated from the vibrating objects and the perceptual estimates of these mechanical properties by listeners. Thus, the psychomechanical studies reported here all deal with the generation of the mechanical information as present in the vibration patterns and the sounds radiated from the moving objects, and the way listeners process this information to derive estimates about physical parameters of these objects.

Pure bouncing occurs when a ball is dropped on a plate from a certain height, loses some energy in the contact between the ball and the plate, leaves the plate in opposite direction with reduced kinetic energy, falls back on the plate again, makes another contact, etc. etc. So, for a purely bouncing ball, the time intervals between the successive moments of contact with the plate follow a relatively simple rule, and the angular velocity of the ball is zero. A purely rolling ball, on the other hand, does not lose contact with the plate, and has a well defined non-zero angular velocity. Various kinds of transitions between bouncing and rolling may occur when an arbitrary object is cast on a plate with an arbitrary speed into an arbitrary direction. First, the object may bounce a few times; if the object is non-spherical, the time intervals between successive bounces can be very irregular and even seem chaotic. After some bounces transitions to rolling may occur, after which the object comes to rest. If an object gains speed, e.g., when rolling down a slope, it may start with pure rolling and, as the speed increases, the object may

lose contact with the plate and transitions to bouncing can occur. The smaller the ball, the sooner such transitions to bouncing will occur. An obvious observation is that listeners can hear these transitions and distinguish between the different types of sound resulting from the different types interaction between a ball and a plate. But much less is known about what kind of information in the sound they use for their judgments, and how this information depends on the mechanical processes that generate the sound. The thesis tries to provide answers to this type of questions by focusing on one specific experimental condition: a wooden or metal ball bouncing on or rolling over a wooden or metal plate.

The research described in this thesis builds on earlier perception oriented work by Houben [44]. In his experiments, the focus was not so much on bouncing but more on purely rolling sounds. Sounds in which transitions from rolling to bouncing were clearly audible were deliberately removed from the sets of stimuli used in the perception experiments. This choice of stimuli was motivated by the fact that the inclusion of bouncing would complicate the processes under study too much. Another aspect of Houben's work was that he established the relation between acoustic features of the rolling sounds, like spectral or temporal variations, and the listeners' perceptual judgments, but he did not attempt to link the acoustic features analytically to the mechanical parameters of the sound producing objects, such as their elasticity, their mass, and their shape. This concentration on perception experiments was again a deliberate choice to focus his work.

As will be seen in the following chapters, the research described in this thesis goes beyond both of these restrictions. Bouncing sounds are included in the stimulus sets, because the presence or absence of bouncing conveys information the listener may use in estimating the size and the speed of the rolling objects. Selecting sounds with as little bouncing as possible may, therefore, remove information that listeners use in making their estimations. Later in this introduction it will be argued that the richness of natural sounds and the redundancy of information present in these sounds may play an essential role in the perceptual processes involved in deriving judgments about mechanical parameters.

As for the experimental methods used, we also include a number of theoretical and experimental analyses of the mechanical events that take place when two solid objects are in contact. These analyses allow to predict how the produced vibrations and sounds depend on the mechanical parameters of the objects and whether different mechanical parameter combinations can lead to similar sounds. By combining these mechanical insights with those from perception experiments it will be shown that for some physical properties, like the size of a ball, listeners

can indeed reliably report their magnitude, while for others, like the plate thickness, they cannot.

In the next sections of this General Introduction we will describe the theoretical and conceptual framework in which the research of this thesis is carried out. This framework will be summarized in a scheme for the various components in the psychomechanical hearing process (Figure 1.1). It consists of a complete research cycle with various levels on the mechanical side, and various levels on the perception side. From the inherent complexity of this framework it will become clear that not all parts of the sketched research cycle can be completed within the limited amount of time of a PhD project. Various parts of the cycle could be addressed successfully. On the perceptual side, e.g., the problem was encountered that the accuracy of the perceptual estimates of the speed of the rolling balls appeared to depend on the experimental paradigm. The theoretical implications will be discussed. Some other gaps in the cycle could only partly be filled in; on the mechanical side, e.g., simulations of bouncing sounds led to satisfactory results for metal balls but not yet for plastic balls. Consequently, a good model could be developed for a metal ball rolling over a metal plate, but for a wooden ball rolling over a wooden plate, the imperfect sphericity leads to temporal variations that are currently not included in simulations. Yet, the sounds of wooden balls are much more pleasant to the ear, and probably more typical of what people expect to hear when listening to a rolling ball. Note that in this process the listener – not a mechanical procedure – determines the quality of the sounds. The theoretical, methodological, and conceptual significance of these findings will be discussed extensively in the course of the thesis.

1.2 Everyday sounds

In everyday life, we gather information about our surroundings via the sound that this environment produces. In the classification of everyday sound-producing events, as introduced by Gaver [24], one of the three branches is labeled “vibrating objects”. Often these objects start vibrating after being in contact with each other. This contact can last shorter or longer, and the type of contact, like hitting, gliding and scraping, forms a continuum of possible interactions. If one of the objects is a ball, the possible interactions are: bouncing, rolling and scraping, and they can also occur in combination, e.g., rolling and bouncing. By studying the perception of these sounds we want to investigate how people use the information in radiated sound that is generated in their surrounding to form an image of

this surrounding. By doing so we investigate an important branch of everyday sound-producing events. The information in this example of a rolling or bouncing ball resides in the geometry of the objects, their materials and their changing positions during the interactions. For example, by listening to its sound we can determine whether a ball is bouncing or not. When imagining a table-tennis ball and a cannon ball falling on the floor, we intuitively know that these events generate quite different sounds. We also know from experience that, when some object is dropped on a carpet, it sounds quite different from the same object when dropped on a wooden floor. These are all examples of information about the mechanical object properties that people can extract from listening to the sound. These are also examples where differences between the mechanical properties of the objects result in different sounds. The discipline that relates the mechanical features of a sound generating system with the auditorily perceived features of this system is called psychomechanics [66][68].

Up to now most psychomechanical research has been carried out with objects that obey relatively simple acoustical descriptions. Examples are the presumed simple relation between the change in the resonance frequency above the water level of a vessel that is being filled [9], or objects that easily and accurately can be described or synthesized with a modal description [63]. Recently, interest has increased in sounds that are more complex, for instance due to damping [68] or due to the variable thickness of a vibrating plate [16]. In this thesis, we will try to go one step further and investigate the perception of mechanical features of a ball rolling over or bouncing on a plate. We will show that many aspects in the contact can provide information to the listener. We already mentioned, the ball might lose contact with the plate while rolling, resulting in bounces, the spectrum of which that can provide extra information about the size of the ball. Not only whether the ball bounces or not, also the bouncing pattern, regular, irregular or even chaotic, may provide information to the listener; as may be the temporal variation in amplitude of the bouncing sounds.

According to the ecological psychology of hearing, the perception of mechanical properties of the ball should be based on coherent information that is obtained via multiple sources of information [26]. If perception, indeed, depends on coherency of the different auditory sources of information, there should be multiple sources of information available to the listener to test this hypothesis. And thus we can only study this process if we use sounds that are naturally rich of such information.

1.3 Information in sound

In many occasions, people seem to trust upon information they obtain from the sounds they hear. For instance, when we close a door we can hear if the door falls into its lock. In the absence of hearing the click, we are informed that the door has not fallen into its lock, and we press the door another time. In this act we seem to trust upon auditory information, because we cannot see the lock that is hidden in the door. We never seem to doubt this information.

In the presence of other sensory information, however, we are often more likely to rely more on that information than on the auditory information. For instance, in vessel filling, people are much more accurate in this task when visual information is also present [9]. This should, however, not lead to the conclusion that information obtained auditorily is inferior to that obtained via vision. The material of a plastic object that perfectly looks like metal can be revealed by tapping on it. Apparently we put, in this case, more trust in the identification of material by sound than we put in its identification by vision.

More anecdotal evidence of the use of auditory information is given by several authors. Most often these examples deal with situations where some object is out of view. Gaver [24] describes the situation where someone is walking in the street and identifies a car behind him by the sound that it generates.

1.4 Psychomechanics

Gaver [23][24] introduced two terms for two different ways of listening, musical listening and everyday listening. In musical listening, we judge sounds on the basis of their musical properties. For instance, we listen to the pitch, the duration and the harmonicity of sounds. In everyday listening, for instance when hearing a resonating sound after hitting an almost empty bottle, we will foremost conclude that it is almost empty. We will hear, but not put any attention to the fact that it is a resonating sound. The scientific problem with this view is that there is, besides such anecdotal arguments, no evidence for such a division in different hearing modes. Therefore, we do not know if there are only two modes of hearing or perhaps even more. Furthermore, we do not know to which extent these two modes differ. In the studies presented in this thesis, we have always asked people to judge mechanical properties of the sound and thus assume that they always operate in the same type of listening, i.e., everyday listening.

Many everyday sounds consist of contact sounds produced by mechanical interaction between solid objects. That part of hearing research that studies the perception of these sounds have already been indicated with psychomechanics. Psychomechanics is the term introduced by McAdams [66][68] to describe the auditory perception of mechanical properties of a sound-producing system. Typically, in psychomechanics a somewhat different approach is taken. Instead of dividing the *hearing process* into two types, the *sound generation* process is divided into two steps, the source or origin of the sound and the path from the source to the listener. The source is often an interaction between two objects, such as hitting a bar with a mallet, a bouncing or rolling ball or a vessel being filled. When someone is asked what he or she hears after being presented with a recording of a violin, he or she is likely to answer that it is a violin, added with details about the work if it is known to the listener. What he or she is not likely to answer is that it is a CD-player or a loudspeaker that he or she heard, although this formally is also a correct answer to the question. This observation indicates that we naturally speak about the content of the sound, and thus about the source that generates the sound, and not about the path the sound takes to reach our ears.

When explicitly asked, we could also judge the quality of the CD-player or the size of the room where the musical piece was played in. We naturally concentrate on the source because this is what gives us the view on the environment surrounding us. In this way we can tune our hearing to different aspects of the same sound, just as in the case Gaver described. In this case, however, there is no artificial split in two ways of listening. Using a split into the source and path of the sound, we can shift the attention of the listener by asking to judge some source properties or to judge some properties of the path towards the listener.

1.5 Acoustics

It is impossible to treat the perception of everyday sound without looking at the whole picture of acoustics. In this research process we want to concentrate on acoustic variations that indicate considerable variation in the percepts associated with these acoustical variations. It would be a waste of time to ask people to judge some features of the mechanical system when we know beforehand that this task is impossible due to lack of variation in the sound corresponding to these features. One hilarious example would be to ask people to judge the color of a rolling ball from its sound. With a thorough analysis of the acoustics of the mechanical system that produces the sound, we can identify the mechanical pa-

rameters whose variation results in large differences in the sound. Although this is, as we will see in this thesis, not sufficient to guarantee that these parameters are identified by the listener, it seems to be a necessary prerequisite.

1.6 Measurement

In the natural sciences, measuring of magnitude involves determining a unique and reciprocal correspondence between some quantifiable property of the magnitude with a standard quantity, called unit, by using an instrument under controlled conditions [85][99]. A measurement is usually distinguished from a count. A measurement is a real number, and is never exact. A count is a natural number and can be exact. For example, we can count the number of balls in a box and be sure of our answer. In other words, there is no measurement error. Measurement is not limited to physical quantities and relations but can be extended to the estimation of magnitudes of any kind. Improvement of measurement techniques has brought the natural sciences to a higher plan. The definition of measurement in the social sciences has a long history. A currently widespread definition, proposed by Stevens [92], is that "measurement is the assignment of numerals to objects or events according to some rule". We should conclude that measurement in psychology and physics are in no sense different. In actual research in physics, the measured quantities are often hidden from the eye and the measurement methods are indirect. Physicists can measure when they can find the operations by which they may meet the necessary criteria; psychologists have but to do the same. They need not worry about the mysterious differences between the meaning of measurement in the two sciences [82]. A metric determines the distances in a space. By creating this metric, not only the measurement points are determined but also the space nearby the measurement points. The easiest way to come from measurement points to a metric is to fit the measurement points with a curve. The distances between points on this curve can now be measured and can serve to define a metric.

These methods of measurement in psychology have, most often, been used on "low level" perceptual phenomena, such as loudness and pitch perception. In this thesis we will use psychometric methods to quantify the human capability on a more cognitive or "higher level" in perception. That is, it will be shown that these methods are also applicable to derive a measure for the listener's capability, not only to detect changes in sounds they hear, but also to interpret and use these sounds to build an image of what is happening around them.

1.7 Psychomechanics research scheme

Both mechanical engineers and psychologists have both studied sound, but each used their own scope of research, mostly without looking to results of the other domain. For the mechanical engineers sound is the vibration of air, while for a psychologist, sound is a percept. In this thesis we will try to link both domains.

The approach combining these domains is schematically illustrated in Figure 1.1, where a sequence is depicted to come from a mechanical system at the left-hand side to the auditorily perceived features of this mechanical system at the right-hand side.

The scheme can be seen as an extension of the strategy for investigating the perception of source events as introduced by Li, Logan and Pastore [60]. They introduced a scheme with three blocks. They labeled the first “*auditory events*” and it corresponds to our block *mechanical system*. The name of this block as used by Li *et al.* suggests that it should be interpreted as being part of the *perceived* reality and thus be placed on the right side of our scheme. It is, however, referred to in the text as describing “acoustic source events”, suggesting it to be part of the physical reality on the left side of our scheme. Also Giordano [28] acknowledged this confounding labeling and used the term “*physical level*” for this block.

The second block was labeled “*acoustic structure*” and corresponds to the combination of our blocks *acoustical radiation* and *auditory information*. The third block of their figure was labeled “*perception*” and corresponds to our block *perceived mechanical features*. We have introduced two new blocks that represent more hidden aspects of the perception process. The *mechanical basis variables* represent the mechanical system, but now using parameters that effectively change the sound in unique ways. A more complete description of the content of these blocks will be given in the discussion of this thesis. Li *et al.* stressed, as we do, that the relation between all blocks should be investigated to understand how the features of the mechanic event are recovered from the acoustic representation of this event. We use this scheme as a guide for the work in the different chapters of this thesis. In the discussion of this thesis we will show which contributions we have made in the investigations of the relations between all blocks, but first we will start by relating the state of the art to this scheme.




FIGURE 1.1: Overview of the psychomechanical hearing process. For details see text.

1.8 Literature

In this section we will review some recent literature and indicate the place of the investigated relations in our scheme.

The left size of the scheme is described by mechanical and acoustic laws found in textbooks on these subjects, such as by Cremer *et al.* [14], Morse and Ingard [71], Junger and Feit [48], Dowling and Williams [18] and many other texts. Note that, however, many aspects of acoustics that are important for sound quality are still being investigated, such as structural damping [12].

McAdams, Chaigne and Roussarie [68] studied the material properties of impacted bars in a psychomechanical study using synthesized sounds. They define three types of variables which they label “mechanical parameters”, “signal descriptors” and “perceptual coordinates”. In terms of our scheme these refer to the *mechanical system*, *acoustic radiation* and *auditory information*, respectively. They conclude that “The results thus validate two aspects of the synthesis model and quantify psychophysically the relations between the mechanical parameters, the potential signal characterization that carry them and the perceptual representation” [68, p. 1319]. They do not address the question what mechanical system is perceived, which is the main question in our research.

A quite similar approach has been taken by Canévet, Habault, Meunier and Demirdjian [10] in their study of the auditory perception of sounds radiated by a vibrating plate. In our terms, they investigated the relation between the same blocks of our scheme as McAdams *et al.* [68]. They also concluded that structural damping parameters are important in the perception of the sound. They found a high correlation between the time decay of the sound and one of the three dimensions of the perceptual space found with a multidimensional scaling (MDS) analysis. The other two dimensions were correlated with the tonal character and the sharpness of the sounds. There was also a correlation between the structural damping and the pleasantness of the sound. Although interesting, sound quality as such is not addressed by our scheme. An important difference between this study and the one by McAdams *et al.* [68] is the type of physical model. Where McAdams *et al.* [68] used a numerical model in which the differential equations are directly discretized, Canévet *et al.* [10] used an analytical model where only the result of the analysis is discretized. The numerical model often leads to more realistic sounds, especially when non-linearities are involved; the analytical model often leads to more insight in the influences of different parameters onto the sound.

In an effort to determine the human capability of determining the size of bouncing balls, Grassi [31] performed a study of the mechanics of a bouncing ball as well as perception tests. He divided the two interacting bodies in two, the non-sounding object and the sounding object. In the case of his mechanical setup, the ball did not radiate acoustically, and hence the plate is the sounding object whereas the ball is the non-sounding object. He does not explain in what respect estimation of properties of the non-sounding object differs from estimation of the properties of the sounding object. Although it seems intuitively more difficult to estimate the properties of the non-sounding object because it is more indirect, this might not be so trivial. In terms of mechanical analysis, the interaction between ball and plate is the source of the vibration. These vibrations are then filtered by the resonating plate and radiated acoustically. Some mechanical parameters, such as the mass of the ball and the density of the plate, have an influence on the interaction; others, such as the plate size, have an influence on the filtering of the vibrations by the plate. A priori we cannot tell if some parameter influencing the interaction leads to larger differences in the sound than a parameter influencing the filtering of the vibrations. Grassi then correlated the results of his perception test with acoustical signal parameters such as RMS power and spectral centroid. The problem here is that these two signal parameters are correlated to a high degree, and he suggested to use synthesized or manipulated sounds to solve this problem. Although Grassi recognized the importance of the contact time and restitution coefficient, he did not use these parameters in the analysis of the results of his perception tests. We will show that these two variables summarize the mechanical system in such a way that we can relate the results of our perception test to these variables.

Kirkwood [50] tested if presentation type has an influence on the perceived length of wooden dowels. He used three presentation types, live generated, binaurally recorded, and monophonically recorded presentations. The performance for the estimation of the length of a rod being dropped on the floor was in these three conditions not statistically different, although there was somewhat more variability in the answers from the data being generated with monophonic presentations. Although Kirkwood warns for using monophonic presentation, from his data we can conclude that we can investigate the perception of mechanical features as well as with binaural or live presentation, at the cost of somewhat more variance in the data. We might presume that the difference will be even less when there are no binaural cues available.

Lutfi [62] investigated the discrimination of intensity in tone complexes. Partial of the tone complexes varied in amplitude in a stochastic manner. The number of partials also varied and it was investigated whether people could make use of the increase in information. The term information is used here, as in information theory, to be synonymous with random variation. In these tests feedback was used to test the limit of the intensity-discrimination capability of listeners. This is a typical example of research that can be placed in a single block of our scheme, that of *auditory information*. No reference is made to the meaning of the sounds, and such research informs us about the capability of the low-level auditory system to detect variations in sound that can be used to infer the meaningful properties of the sound in the higher-level blocks of the scheme.

Another example of an approach that concentrates on the right-hand side of the scheme is given by Houben and colleagues [40][42][43]. In one study, Houben, Kohlrausch and Hermes [42] showed that listeners are more or less capable of determining the speed and size of a rolling ball. They tried to find the cues by which the listeners are capable to do so, by using a computational model for the auditory system. Two features were investigated in their study; the centroid of the specific loudness and the auditory roughness. The problem with these cues is that they both vary in a similar manner with the speed and the size of the ball, and so cannot be used to judge these mechanical parameters independently. In another approach by the same authors [43] they tried to generate new rolling sounds by combining the spectral content of one rolling sound with the temporal variations of another rolling sound. They again asked listeners to judge the size and the speed of the ball. In this way they could determine whether temporal or spectral properties were responsible for the speed and size judgments of the listeners, and it turned out that for these sounds the spectral content dominated in both cases, while temporal variations were shown to play only a minor role. In yet another study, Houben [40] added amplitude modulation to the sound. The reason for doing so, was that in the rolling sound of wooden balls there is always some amplitude modulation present, which provides information about the size and speed of the ball. This information could be used to disambiguate the other cues that give similar information about the size and speed of the ball. Listeners indeed varied their answers according to these added amplitude modulations, and Houben concluded that when judging the speed of a rolling ball, the participants' judgments correlate predominantly with the angular speed corresponding to this amplitude modulation. If we project their methods on our scheme, we can say that they built a computational model for the right-hand side of the scheme,

to overcome the problem that many of the variables in this domain cannot be measured directly.

In contrast to the mentioned work, we will try to address all blocks of our scheme and we will explain some of the more hidden blocks, the *auditory cues* and the *mechanical basis variables*.

1.9 Goals and main contributions of this thesis

The goal of this project was to investigate what kind of auditory information listeners use in judging the properties of a mechanical system, how well this auditory information represents the parameters of this mechanical system, and how accurately a listener can use this information to perceive the physical properties of the mechanical system, such as the speed, the size, or the material of the objects the system consists of. Our object of study was that of a ball bouncing on or rolling over a plate.

In order to realize this goal, a novel research paradigm has been described to investigate the auditory perception of mechanical features of a sound generating system. It contains elements from the paradigms suggested by Gaver [23][24] and Li, Logan and Pastore [60]. In essence, it is an extended analysis/resynthesis paradigm but, both on the mechanical and on the perceptual side, the system is split up in various levels of description. Various elements of this novel paradigm are well known: some originate from psychomechanics, some from signal detection theory, and others from psychophysics. They are used, however, in a new configuration that systematically distinguishes three levels on the mechanical side and three levels on the perceptual side. By studying the relation among the levels and the flow of information between them, we expect to reveal the processes underlying the perception of mechanical features of a sound generating system.

The current thesis extends the goals of the thesis of Houben [40]. He determined how well listeners can judge the size and the speed of a rolling ball and investigated on which kinds of *signal properties* the listeners' judgments are based. This thesis goes a step further and, as mentioned, aims at finding out what *mechanical parameters* underlie the acoustic information used by listeners. To this end, the current method concentrates on the flow of information from the sound generating system to the image the listener reconstructs from auditory information coming from this system. By addressing separate points in the processing

line that is formed by this information flow, we will describe the individual processing steps of this flow as studied in this PhD project.

The psychomechanical research paradigm introduced in section 1.7 is the theoretical and methodological contribution of this thesis to psychomechanical research. Essentially, this psychomechanical approach requires sound mechanical descriptions of the system under study, preferably at the three levels described for the mechanical side of this research scheme. It became apparent that the existing models of bouncing and rolling were too simple to explain some phenomena that we observed and wanted to investigate for their role in the perception of mechanical features of these interactions. Hence, this thesis contains contributions on three mechanical phenomena, and on three perceptual phenomena:

- First, it presents a model of bouncing that builds upon the mechanical point impedance of a plate, the contact stiffness, and the contact time. This model explains the effect of various mechanical parameters upon the spectrum and temporal pattern of the generated sound.
- Second, an interference pattern is described that appears in the recordings of the rolling sounds. It will be shown that this pattern arises from interference between the vibrations induced at the point of contact between ball and plate, and the reflections of these vibrations at the edge of the plate.
- Third, a time-domain model and simulations of a ball rolling over or bouncing on a plate is presented. Using this model we could simulate four different types of contact between the ball and the plate, sometimes regular and sometimes irregular or chaotic.
- Fourth, the relevance of the restitution coefficient, and the consequent temporal pattern, on the perception of the size of the bouncing ball is studied. Remarkably, it appeared that this pattern plays only a minor role in the size judgments of the listener. The spectral properties of the impact sound appears to be more important.
- Fifth, the flow of information between auditory information and perceived mechanical features (see Figure 1.1) was studied for the acoustic interference patterns resulting from edge reflections. Listeners can in general distinguish between the sound of a ball rolling towards the edge of a plate and that of a ball rolling away from it. They are, however, not able to use this information in judging the rolling direction. This spectro-temporal acoustic pattern is thus an example of auditory information that does not support

the perception of mechanical features. In a similar way, the perception experiments involving different plate thicknesses revealed that the related auditory information is perceived by the listeners, but cannot be used to judge the plate thickness.

- Sixth, for the perceived mechanical features of size and speed of a rolling ball we gained new insights into the relation between the three processing levels on the perceptual side of our psychomechanical research scheme. Size perception is most robust, because it is equally accurate in an absolute estimation paradigm and in a paired comparison paradigm. In contrast, the accuracy of speed perception diminishes when this parameter has to be judged in isolation, compared to the accuracy in a paired comparison experiment. Taken together with the fifth item on this list, these results suggest that the auditory perception of mechanical object properties is a layered process, and that detection and estimation take place at different levels of such a layered process.

1.10 Outline

This thesis is organized in such a way that specific aspects of the ball-plate interaction are first investigated in terms of mechanics, followed by a perceptual study of the auditory consequences of the found mechanical relations. We will do this for bouncing behavior, for an interference pattern arising from reflections at the edge of the plate, and finally for rolling behavior.

In Chapter 2, the mechanics of a ball bouncing on a plate are investigated. We will show that the temporal pattern of the bouncing process can vary, among others, with the size, the material and the height from which the ball is dropped, and we will try to describe this system in terms of mechanical basis parameters. In Chapter 3 we will relate these findings to the perception of the ball size from bouncing sounds. By manipulating the temporal pattern of the bouncing sounds, we will show that listeners do not attribute much weight to this temporal pattern. The frequency content of the bouncing sounds is given more weight in judging the size of the bouncing ball.

When a ball rolls over a plate, it generates vibrations in the plate that are reflected at the border of the plate. This results in a spectro-temporal pattern in the sound recordings. This pattern, described in Chapter 4, gives information as to whether the ball rolls towards the edge of the plate, in which case the spectral peaks increase in frequency, or away from the plate edge, in which case the spec-

tral peaks decrease in frequency. In Chapter 5, we will show that listeners are well capable of distinguishing between sounds with such a decreasing or increasing interference pattern but do not use this information in judging whether the ball rolls towards the edge of the plate or away from it.

In the remaining chapters, we will focus on rolling sounds as a whole. In Chapter 6 we will describe a numerical model that is able to generate rolling sounds, together with measurements and simulation results. It will turn out that, for a large class of rolling sounds, a relatively realistic description can be given. This applies in particular to small metal balls rolling over a wooden or metal plate when the sound is not too monotonous, for instance because the balls are bouncing somewhat. In the two next chapters, we will look into the perception of the rolling sound, using psychometric functions in Chapter 7, and using multidimensional scaling (MDS) in Chapter 8. From the results of these experiments we will quantify the human capability to auditorily perceive the ball speed, ball size and plate thickness. As for plate thickness, listeners are well capable of hearing the differences but, as in the case of the interference pattern, they do not use this information in judgments on plate thickness. For speed judgments listeners are more accurate when comparing two rolling sounds one after the other than when asked to judge the speed of the ball based on one sound. Regarding size judgments, the two methods lead to equal results.

A theoretical framework will be presented that allows to organize the processes underlying the perception of mechanical features of a sound generating system. This is summarized and discussed in the final chapter of this thesis.

2 Bouncing of a Ball on a Plate: A Model and Measurements Based Upon the Mechanical Driving-Point Impedance of the Plate

Abstract

The restitution coefficient of a ball bouncing on a plate can be calculated from Hertz' contact theory, Newton's second law and the response of the plate. We have adapted this method by calculating the plate response from the point impedance of the plate. This is not only a more comprehensible approach than the classical analysis [111], but the results are also more generally applicable. For instance we will show that, by altering the plate impedance, this new model can also be used to describe bouncing at the edge of the plate. In order to evaluate model predictions, we have set up the following experiments. First we measured the plate impedance to verify its calculation from the mechanical plate properties. We then used the model to calculate the influence of the plate thickness and we calculated the restitution coefficient both in the middle of the plate and at the edges, where the restitution coefficient is much lower. Measurements showed good resemblance between predicted and measured restitution coefficients, as long as the ball was hard compared to the plate. In the opposite case, where the ball was soft compared to the plate, the model was not applicable. Measurements show a completely different behaviour of the restitution coefficient in this case.

2.1 Introduction

When a ball is dropped onto a plate it will bounce some time before it comes to rest. From practical experience we know or can imagine some situations where a ball bounces very long or very short. For instance, a ping-pong ball or a rub-

ber bouncing ball remains bouncing for a very long time while a cannon ball on grass, a ball made of deformable clay or the silly putty ball used as example by Cross [15] are examples of balls that bounce only a few times, if at all.

The parameter used to describe this difference in behavior in bouncing is the restitution coefficient, which is defined as the ratio of the velocities shortly before and after the bounce. The reduction in speed is a result of the ball losing energy in each bounce. For a linear process we would expect that this loss of energy in each bounce is a constant fraction of the kinetic energy of the ball, and the restitution coefficient would be constant throughout the bouncing process. Due to non-linearities this is, however, not always the case.

This restitution coefficient can take values between 0 and 1. The practical meaning of a high restitution coefficient, close to one, is that the ball will keep on bouncing for a relatively long time whereas in the case of a low restitution coefficient, close to zero, it will come to rest on the plate after a few bounces.

One very remarkable observation in the bouncing of a ball on a plate is that, for some balls and some plates, the ball bounces very well on the middle while at the corner of the plate it remains lying on the table after the first impact, without bouncing at all. In this chapter we will derive a physical model that describes this effect. One specific motivation to model this phenomenon is our interest in determining which information about the bouncing events can be extracted by just listening to its sound. The restitution coefficient is potentially a good cue to be used by human listeners to determine the size of a bouncing ball but it has not been studied in other perceptive studies of bouncing, for instance those by Grassi [31][84].

We will analyze two setups that result in different bouncing behavior of the ball. In the first setup we used metal balls bouncing on MDF plates. In a second setup we used plastic balls bouncing on metal plates.

The mechanics of bouncing balls have been studied from a number of perspectives. The first important contribution was made by H. Hertz in 1881 [37]. He studied the contact stress of two spheres. His results are well known and applied in many areas. In a publication in 1941, Zener [111] used bouncing to study the reaction of plates to forces of short duration, and established the influence of many material properties, including geometrical ones, on the restitution coefficient. More recently the calculation of the restitution coefficient, including the influence of position of the plate, was studied by Sondergaard *et al.* [90] whose interest originated in a completely different area, namely gas dynamics.

Our interest comes from studying auditory perception of object properties such as the ball size. People can determine the ball size by listening to the sound of a bouncing ball, as was shown by Grassi [31], and we will investigate further in the next chapter. But they can also discriminate between balls of different sizes by listening to the sound of rolling balls [40]. Also for rolling sounds the vibration originates from the contact between ball and plate. In order to better understand the mechanical behavior of the contact, we try to calculate first the case for a bouncing ball, because here we can do some well-defined measurements and there are some published experimental results available. The model parameters leading to correct bouncing ball predictions can be used as a starting point for the model for rolling sounds.

In the case of bouncing the contact is, however, generally so short that, when the balls bounces sufficiently far from the border, the vibrational waves reflected at this border of the plate do not return to the point of impact during the time of impact. During the time of impact the plate can thus be considered as infinitely large. The point of impact does, in an infinitely large plate, only move during the time of impact. A vibrational wave travels from this point outward. When the plate is finite in size, this wave will be reflected at the edges of the plate, and the plate continues to vibrate. When the vibrational waves return at the point of impact during the contact of a single bounce, this will be of influence on the restitution coefficient. This happens when the ball bounces close to the edge of the plate. The critical distance from the border up to which the plate can be considered infinite during the contact between the ball and the plate will be studied later in this chapter. Furthermore, in our experiments, the plate comes to rest in between the bounces, except only for the last few bounces, simplifying the analysis further.

There are several reasons why bouncing of balls is interesting for the study of rolling sounds. First of all, it is much simpler to compare the numerical and experimental results for bouncing than for rolling. Therefore, any parameter involved in rolling as well as in bouncing can best be measured using the latter, an approach which was also applied by Brilliantov and Pöschel [8]. Furthermore, bouncing is often included in rolling.

Because people have experience with all kinds of bouncing objects, we would like to have a somewhat broader view on this bouncing process, compared to the mentioned studies where a single setup was studied. We will show different bouncing behavior for two setups, which is of importance to understand what people may consider natural bouncing behavior and what not. If we want to test

the perceived relation between the restitution coefficient and the ball size, as we will do in the next chapter, we should know what this relation is in the physical world and under which conditions it is present.

2.2 Theoretical analysis

Experimental data comparing the bouncing of a 10 mm steel ball on a steel plate or on a granite plate were published by Stensgaard and Lægsgaard [91]. For the granite plate the restitution coefficient seems constant; for the steel plate it increased with decreasing impact velocity. Such data can provide information about the energy dissipation causing a reduced restitution coefficient.

The loss of energy can be due to many effects. For instance:

1. Viscoelastic deformation of the ball [15],
2. Elasticity of the ball (it remains vibrating after the contact),
3. Viscoelastic loss in the contact region [55][20],
4. Elastic waves in the plate [111][29],
5. Plastic deformation of ball or plate.

Measuring the restitution coefficient for low impact velocities, which could provide answers regarding the correct model for the viscoelasticity, has been shown to be difficult [20]. This is partly due to the fact that the gravitation influences the way the ball comes to rest. We neglect the viscoelasticity in our analytical analysis. In the numerical simulations it is possible to include these effects.

When the ball bounces on a plate, the primary loss of energy happens during the contact with the plate. During the contact-free phase the motion is governed by the gravitational forces and inertia of the mass of the ball. During the contact the gravitational forces can be neglected. Due to the very short time in which the ball reverses its direction, the inertia of the ball puts a much larger force on the plate. Because ball and plate are not perfectly rigid, the ball leaves the plate with a slightly smaller velocity. Due to this effect, the restitution coefficient varied in our experiments from about 0.95 to 0.5.

Throughout the analysis this chapter, a Kirchhoff-Love model for the plate and a Hertzian model for the plate are used. These models have their intrinsic limitations. The Kirchhoff-Love model is a thin-plate approximation because it

neglects the effects of rotatory inertia and shear [69]. Especially in the thickest of our plates these effects may play a role. The Hertzian model assumes that the strains are small, and no plastic deformation occurs [46]. These limitations might contribute to the discrepancies between the model and the measurements later in this chapter.

2.2.1 Energy balance

At the moment just before impact, the following kinetic energy E_z is stored in the ball:

$$E_z = 1/2 m_s v^2. \quad (2.1)$$

The energy difference due to the different velocities of the ball, before and after the impact, is transferred into the plate. This difference in velocity, of course, also determines the restitution coefficient. We can also calculate this energy from the resulting plate vibrations. Starting from the plate equation, see for instance [30],

$$\rho h \frac{\partial^2 W}{\partial t^2} = \frac{\partial^2 M_x}{\partial x^2} + \frac{\partial^2 M_y}{\partial y^2} + 2 \frac{\partial^2 M_{xy}}{\partial x \partial y} + f_z(x, y, t), \quad (2.2)$$

with, as usual, W the plate displacement, ρ the plate density per area, h the plate thickness, x, y the coordinates along the plate and M_x, M_y, M_{xy} the moments on the plate, and finally $f_z(x, y, t)$ the external force applied by the bouncing ball. We can calculate the part of the energy in the plate after the impact, by calculating the force,

$$F_z = \iint f_z dS, \quad (2.3)$$

from the displacement and the moments. The capital letter F denotes the complete force on the whole of the plate, and S is the surface of the plate. We can calculate the energy from the force

$$\begin{aligned} E &= \int F dW \\ &= \int F \frac{\partial W}{\partial t} dt. \end{aligned} \quad (2.4)$$

Because all terms in Equation (2.2) are essentially forces, we can use Equations (2.3) and (2.4) to get

$$\rho h \iiint \frac{\partial^2 W}{\partial t^2} \frac{\partial W}{\partial t} dt dS - \iiint \left(\frac{\partial^2 M_x}{\partial x^2} + \frac{\partial^2 M_y}{\partial y^2} + 2 \frac{\partial^2 M_{xy}}{\partial x \partial y} \right) \frac{\partial W}{\partial t} dt dS = E_z. \quad (2.5)$$

This calculation is easily done in a numerical domain where the plate vibrations are known with great precision for each moment and each position in space. An example of a model that enables us to do such an analysis will be discussed in Chapter 4.

In this chapter we will describe an analytical model, for which the plate vibrations do not need to be known. This can be achieved by approximating the real, finite-sized plate with an infinite one. The behavior of the plate is then sufficiently described by the mechanical plate impedance. Later we will extend this model to be valid in the proximity of one plate edge, that is to a semi-infinite plate.

2.2.2 The contact between ball and plate

When a ball is pressed against a plate, the deformation of the ball and plate has the effect that the center of the ball comes closer to the horizontal midplane of the plate than in the undeformed condition. This effect will be indicated with α :

$$\alpha = u_p - u_s, \quad (2.6)$$

where the displacement of the plate is labeled u_p and that of the sphere u_s . The reference position $\alpha = 0$ is defined for the equilibrium in which the ball is in contact with the plate without putting any force on it. When a force is put onto the ball, α becomes negative. The remainder of our analysis is only valid when the ball and plate are in contact, thus, α can never be positive. The quantity α was called “approach” by Zener [111] and others, but we will not use this term to avoid confusion.

The deformation of the plate can be split up in two types: First, the local deformation at the point of impact. Because the area of deformation is very small compared to the thickness of the plate, we can approximate the plate with an elastic half space and use Hertz’ theory to calculate the resulting force of this deformation,

$$F = \kappa \alpha^{3/2}, \quad (2.7)$$

where

$$\kappa = \frac{4}{3} \sqrt{R} \left(\frac{1 - \mu_p^2}{E_p} + \frac{1 - \mu_s^2}{E_s} \right)^{-1} \quad (2.8)$$

depends on the radius of the ball, R , and the stiffness of the ball and the plate material represented by E and μ , which are, respectively, Young's modulus and the Poisson ratio of the sphere and plate, indicated again with subscripts s and p

The second type of deformation is due to the deformation of the plate as a whole by the force of the ball which is transmitted via the contact area. In this case we use the thin-plate approximation, to calculate the mechanical point admittance. β ¹ This is exactly the reciprocal of the mechanical point impedance of a plate, which has been calculated by Cremer [14] and also Junger and Feit [48], thus

$$\beta = \frac{1}{4h^2} \sqrt{\frac{3(1 - \mu^2)}{E\rho}}. \quad (2.9)$$

Their method does, however, not apply for the border of the plate, but for this special case the admittance was calculated by Eichler to be 3.5 times the original admittance [19].

We can use β as an ordinary admittance,

$$v_z = F\beta. \quad (2.10)$$

The source of the force lies in Newton's second law, which reads in this case:

$$-F = m_s \frac{d^2 u_s}{dt^2} \quad (2.11)$$

where m_s is the mass of the sphere. After differentiating Equation (2.6) twice with respect to time, substituting Equations (2.7), (2.10), (2.11) yields

$$\frac{d^2 \alpha}{dt^2} = \kappa \alpha^{1/2} \left(\frac{3}{2} \beta \frac{d\alpha}{dt} + \frac{\alpha}{m_s} \right). \quad (2.12)$$

Equation (2.12) was solved numerically to determine the outgoing velocity of the release of the ball for a given impact velocity. The velocity is, throughout this

¹In the original article, Zener [111] uses, instead of the plate impedance, a derivation where the plate size is made infinitely large after using a modal approximation. Modes do, however, not exist on infinitely large plates and therefore this derivation is not fully comprehensible. It does however, lead to exactly the same result.

calculation, strictly perpendicular to the plate and when the outgoing velocity is divided by the incoming velocity, in absolute sense, we obtain the restitution coefficient.

It is possible to find an analytic expression for the time the ball and plate are in contact during the bounce, which we will call the contact time for short (see Chaigne and Doutaut [11] and [20])

$$t_c = 3.22 \left(\frac{m_s}{K} \right)^{2/5} v_{in}^{-1/5}, \quad (2.13)$$

and the dimensionless inelasticity parameter as introduced by Zener [111]:

$$\lambda = \beta \frac{m_s}{t_c}. \quad (2.14)$$

The restitution coefficient is a function of λ only, but, due to the non-linear nature of Equation (2.13), there is no analytic expression, that we know of, to solve it. The relation has been plotted graphically by Zener [111]. As we will see in the next section describing measurements, this calculation is accurate as long as the main loss of energy is through the plate that is set into vibration.

2.3 Measurements

Our derivation of the equations which, after numerical simulation, lead to the restitution coefficient, depends on the mechanical admittance of the plate. According to the theory, an infinite plate has a real impedance, which means it cannot remain moving when the ball has left the plate. Of course, due to its finite size, and thus inevitable reflections, the plate remains vibrating for a while after the hit. Because our calculations of bouncing do not extend to the time after the impact, we are really interested in the impedance the given plate would have if the plate was infinite in size. The impedance is independent of frequency, therefore we do not calculate the impedance in the frequency domain as is usually done.

In our first setup, a large difference exists in stiffness between the material of the ball (steel, $E = 215$ GPa) and that of the plate (wood $E \approx 5$ GPa). Due to this effect, it is mostly the plate that deforms during the bounces. In our second setup, the situation is about reversed. The plastic balls ($E \approx 0.5 - 2$ GPa) are very soft compared to the metal plates ($E = 70$ or 110 GPa).

2.3.1 Experimental setups

We used two different measurement setups to measure the restitution coefficient of bouncing balls.

The first measurement setup consisted of a plate of MDF (Medium Density Fiberboard, wood particles that are compressed and glued together), with a thickness of 6, 12 or 18 mm, and a size of 50x120 cm. The plate rested on four air balloons, to minimize the effect of the suspension. Various balls were used ranging from 10 to 19 mm in diameter and made of steel. The balls were held in position by an under pressure in a cone-shaped device and were dropped by letting air in via an opening at the top of this device. The vibrations of the plate were measured with accelerometers, placed close to the point of impact. These were connected to their special power supply and pre-amplifier and from there the signal was fed into a high quality PC sound-card.

For the second measurement setup, we used aluminium or brass plates of size 100x20 cm and an average thickness of 10 mm for the aluminium plate and of 5 mm for brass plate. The plastic balls had diameters of 10 or 20 mm and were made from polypropylene or teflon, two materials differing in density and Young modulus. Also aluminium bars were used to bounce the balls on, but the results were equal to the metal or brass plates.

2.3.2 Impedance

We measured the impedance by hitting the plate with an impact hammer at different positions of the plate. We obtained the speed of the plate by integrating the acceleration which was measured with the accelerometer. The contact force was obtained from the hammer. Then the impedance was obtained by simply dividing the speed by the force on the highest point of the force curve. The results are found in Table 2.1.

The theoretical values given in the table are based on Equation (2.10) where the Young's modulus, E , was determined by measuring the bending of a small, simply supported bar of the same material. Comparing the theoretical values with the measured values, it appears that the measured values are 5 to 25% lower.

2.3.3 Case 1: Bouncing of a hard ball on a soft plate

We measured the restitution coefficient for three different plates. The three plates differ in Young's modulus as well as in thickness. In our model we only need to

TABLE 2.1: Mechanical driving-point impedance of the various positions of the plate as a function of the plate thickness. Three values are given for the middle of the plate: the first is in the geometrical center of the plate, the others are 10 and 20 cm away from that center. Driving-point impedances are given in Ns/m

h	<i>theory</i>	<i>middle</i>			<i>edge</i>		<i>corner</i>	
	Z_m	Z_{m1}	Z_{m2}	Z_{m3}	Z_{m4}	$\frac{Z_{m1}}{Z_{m4}}$	Z_{m5}	$\frac{Z_{m1}}{Z_{m5}}$
18 mm	1692	1271	1607	1313	289	3.28	169	7.57
12 mm	724	687	627	667	186	3.65	92	7.37
6 mm	181	147	147	155	49	3.00	30	4.90

change the admittance of the plate. The results calculated numerically with the model are compared to the measured data in Figure 2.1. The model and measurements show an increased restitution coefficient for lower impact velocities. Furthermore both the model and measurements show an increased restitution coefficient for thicker plates. They differ in that the model overestimates the effect of the plate thickness. For the thicker plates the predicted restitution coefficient is too high, but for the thinnest plate the measured restitution coefficient is higher than the one predicted by the model. Overall, there is a reasonable agreement between the measured and calculated data. The error is about 10 percent, which is roughly equal to measurements reported elsewhere [90].

If the restitution coefficient predicted by the model is higher than the measured value, this might be explained by loss of energy not taken into account by our model. It is thus surprising to see that for the thinnest plate the predicted restitution coefficient is actually lower than the measured value. Therefore it is more likely that this discrepancy is due to the intrinsic limitations of the Hertz model. Possibly for the lowest plate the deformation depth is not small compared to the plate thickness.

Another effect predicted by the model is that the gradual decrease of bounce height results in elongation of the contact between ball and plate during the bounces. This causes the sound of the bounces not only to become softer but also somewhat duller or lower in frequency.

2.3.4 Boundary influence

The restitution coefficient is substantially lower for a ball bouncing near the border of a plate. In order to be able to model this behavior we need to replace the infinite plate with a finite one. A semi-infinite plate model, with only one, straight edge was derived by Eigler [19]. He found an admittance at the edge which is 3.5




FIGURE 2.1: Comparison between theory and measurements for the restitution coefficients of a 10 mm steel ball on three different plates. The different plate thicknesses are indicated above the panels. Bounces from three bouncing series, connected with a dashed line, are compared with theoretical values, plotted in a continuous line.

times larger than for the infinite plate. The admittance has also a small imaginary part, but it is only 1.5 percent of the real part and therefore we will neglect it. This 3.5 times difference in admittance is valid as long as there are no other edges close to the point of impact, thus not too close to the corners and the opposite edge. In our experimental setup this condition is fulfilled for bouncing more than a few centimeters away from the corners.

By using Eigler's plate-edge impedance instead of the infinite-plate impedance, we can use our model to calculate the restitution coefficient of a ball bouncing near the edge of a plate. To test this extension of the model, we measured the restitution coefficient of a ball bouncing in the middle of the plate, near the edge and near the corner of the plate. The results are shown in Figure 2.2, and compared with the theoretical values using the infinite plate impedance and the semi-infinite plate impedance for, respectively, the calculation of the restitution coefficient at the middle and at the side of the plate. The restitution coefficient is somewhat lower than the theoretical values, as noted before. Of course it is not possible to bounce exactly on the edge of the plate, so we should extrapolate the values approaching that border. This leads to a value of somewhat less than 0.6, which is very close to the theoretical value, but the difference in the restitution coefficient between a ball bouncing on the middle of the plate or near the edge of the plate is lower for the measurement than for the predicted values. The figure




FIGURE 2.2: The restitution coefficient as a function of the distance to the border. The two indicated values are the theoretical values for an infinite plate, 0.87, and at the plate edge, 0.63. The distance indicated on the abscissa is the closest distance from the impact point to the plate edge. The ball was 10 mm in diameter, made of steel and hit the 18 mm thick MDF plate with a velocity of 1.4 m/s.

also shows that there is, as expected, no difference between the long and short side of the plate, as long as we do not approach the corner of the plate.

2.3.5 Case 2: Bouncing of a soft ball on a hard plate

So far, we assumed that the reduced kinetic energy of the ball during a bounce is put into the vibration of the plate. For our second setup, this is not the case. Because the ball is soft compared to the plate, the ball starts vibrating. We measured these vibrations by putting an accelerometer on top of the ball while it bounced on the plate. The result is plotted in Figure 2.3. The same measurement with a steel ball does not lead to much vibration of the ball after contact. Note that the accelerometer influences the vibrations of the ball, because of its presence. Therefore it is hard to estimate how long the ball remains vibrating in the case the accelerometer is not present. We have not been able to create a theoretical model, based on a vibrational analysis of the system, explaining the experimental results of this system, as shown in the next section. Therefore our explanation that the ball vibrations are the main cause of the observed behavior of the restitution coefficient is so far only a hypothesis.




FIGURE 2.3: Vibration of the ball during a bounce. During the period between about 0.05 and 0.18 ms the ball and the plate are in contact. After this first half period of a sine, the contact is lost, but the ball remains vibrating.

A simple analytical model for this system is not known to us. Due to the complicated vibrations of the ball, we could not simplify the vibration of the ball as was possible with the vibration of the plate in the case of the hard ball bouncing on the soft plate. The experimental results from this system are, however, astonishingly simple.

To generate the plots shown in Figure 2.4, we dropped balls on the plate without special apparatus. From the time in between two bounces, the velocity of impact can be calculated, and from the ratio of two consecutive bounces, the restitution coefficient was calculated. The balls were dropped ten times to obtain more data points. The plate vibrations indicate the moments of impact, and these can be measured with an accelerometer. Each circle in Figure 2.4 represents one bounce. The first bounce is not shown because the incoming velocity is unknown. For each plot, ten bouncing sequences were used. For the lower impact velocities, at the end of the bouncing process and shown on the left in side of the plots of Figure 2.4, the plate does not come to rest inbetween the bounces. These vibrations influence the bouncing, and therefore the plots become more scattered for lower impact velocities. This hypothesis was verified by putting the plate on a damping foam layer. This indeed reduced the scattering of the datapoints for low velocities.

For a polypropylene ball, shown in the top two panels of Figure 2.4, the restitution coefficient is independent of the impact velocity. It is also independent of the ball size, always about 0.8. A ball made of a different material, teflon, also resulted in the same restitution coefficient of 0.8. The density of teflon is 2.4 times higher than the density of polypropylene and the Young modulus of teflon is 1.86

times lower than that of polypropylene. Probably these two compensate each other.

Some balls of other materials were also used to measure the restitution coefficient. Some plastic balls had a velocity-independent, constant restitution coefficient of 0.9. Although the material of this ball is not known for certain, and we have not measured the Young modulus for this ball, it is clear that the restitution coefficient can be different from 0.8 and depends on the material of the ball.

A result not shown in the figures was that also the plate geometry had no influence on the restitution coefficient. All plates and even bars we tried resulted in the same restitution coefficient. This indicates that the plate can be considered as an infinite half-space during the contact, if the plate is sufficiently thick.

The bouncing behavior observed here is different from what we observed in case 1, the restitution coefficient is independent of the impact velocity in case 2 where it depended on the impact velocity in case 1. To verify that this different behavior compared with our other measurement setup resulted from the ball being very soft, and did not result from other differences in the measurement setup we dropped metal balls on the plate. The results were very similar to our measurement setup where a metal ball bounced on a wooden plate, the restitution coefficient depends on the impact velocity and varies with the ball size. In this case, however, the plate is not soft compared to the ball and this should be taken into account when calculating the contact stiffness coefficient κ , for the rest the analysis of the wooden plate setup would be valid for this setup.

2.4 Discussion and conclusion

We have adapted the long known model for a sphere bouncing on a plate to be based on the driving-point impedance of the plate. We have verified this model step by step by measuring the plate impedance, contact time and the influence of the plate thickness and the finite size of the plate. We consider the derivation of the model to be more comprehensible than the usual approach, since we do not use the eigenfunctions of the plate, which do not exist for the infinitely large plate. Also, to examine the various effects we can use the separately measurable plate impedance instead of the dimensionless equations of Zener [111]. Generally the measurements and calculations for the restitution coefficient differed by 10 percent or less. There were two exceptions. First, a larger difference was observed the restitution coefficient of the thinnest of our plates. Second, for a soft ball bouncing on a hard plate we have no theoretical model.




FIGURE 2.4: Restitution coefficient as function of the impact velocity of the ball. Note that, in a natural bouncing process, the first bounce has the highest impact velocity and is thus plotted on the right side of the graphs. The six panels correspond with different setups, that is, plate and ball materials or geometries. For a discussion of the different results of these setups, see text. Note the differences in x-axis range.

The theory developed here has many implications for the perception of bouncing balls that have so far not been addressed. There are two obvious sources of information that aid the perception of the size of a bouncing ball. The first is the spectrum of the impact sound. Intuitively we know that a large ball generates more low frequencies than a small one, when bouncing on a plate. The parameter that determines the lowest frequency is the contact time. This contact time is a function of the size and weight of the sphere, as well as the elasticity parameters of the ball and plate, but neither of the geometry nor the weight of the plate. Of course, the vibrations pass through the plate before being radiated, and thus the plate geometry has a large effect on the total spectrum of the sound of a bouncing ball.

The other source of information is the restitution coefficient. Large, heavy balls have a lower restitution coefficient than small, light balls. One interesting aspect about this restitution coefficient is that it is very robust against other influences. These influences can consist of background noise or spectral changes of the sound, for instance when a recorded version is played back with a very low quality. The restitution coefficient is robust in the sense that it can provide information to the listener, as long as the time of the impact is detectable. An interesting case is the difference between the restitution coefficient as well as the spectral properties of the sound of a ball bouncing in the middle or at the edges of the plate. Perception experiments will have to reveal whether the listener will correctly identify this effect or will add more weight to the spectral or temporal sources of information. We have seen examples of mechanical setups that did not change while the restitution coefficient did, namely a metal ball bouncing near the edge of a wooden plate. We have also seen examples of mechanical setups that did change while the restitution coefficient did not, namely the plastic ball bouncing on a metal plate. Therefore the restitution coefficient seems an unreliable source of information about the mechanical parameters of the sound-producing process that cannot be trusted by the listener. The same can, however, be said about the spectral information, that changes for instance due to room acoustics.

A situation that resembles rolling more closely is the bouncing of a ball on a vibrating surface. It appears to be a classical example of chaos [34] although this is contradicted by others [61] [17]. Either way it seems a good possibility to study the (quasi-)chaotic behavior without including the large plate model. Typically, in the mentioned articles, the bounce is idealized. The contact time is assumed to be zero. There is no nonlinear contact, that is, no Hertz or other non-linear-

contact model, and the restitution coefficient is a constant, not depending on the incoming velocity.

The mechanics of a bouncing ball will help us to develop a numerical model for rolling sounds. As we will see in Chapter 4, the experimental results shown in this chapter can be used as a test case to benchmark the numerical model. We will compare these experimental results with results from simulations with a numerical model. Furthermore the restitution coefficient for very low velocities can indicate the correct damping model for the contact, as explored by Falcon *et al.* [20].

3 | The Bouncing of Balls: Influences of Spectral and Temporal Variations on Perception¹

Abstract

In this chapter, the perception of the size of a bouncing ball based only on the generated sound is studied. We will start by identifying two parameters in the mechanical analysis of bouncing balls that control the spectral and temporal character, respectively, of the generated sound. A signal-processing method to modify the spectral and temporal properties of bouncing ball sounds is presented. In a perception experiment we found that at least some of the participants were able to recognize the natural relation between the spectral and temporal properties of these sounds. In a second experiment it was tested if listeners relied more on the spectral or on the temporal properties in their judgment of the size of the bouncing ball. The spectral cues appear to be used in a more consistent way and, on average, have more influence on the participants' judgment.

3.1 Introduction

In rolling sounds, spectral and temporal effects are not easily separated mathematically [43]. To investigate their individual contributions to auditory perception, we chose to investigate a different kind of contact sound generated by a ball on a plate, i.e., the sound of a ball bouncing on a plate. We will show that for this mechanical interaction the spectral and temporal information can be manipulated independently.

Most perceptual research of impact sounds has concentrated on the impacted plate or object, for instance [63][54][68][56]. As far as bouncing is concerned,

¹Part of this work was presented at the CFADAGA04 Conference [96]

Grassi [31] has shown that people can detect the size of a bouncing ball with remarkable precision. By separating temporal and spectral information we will investigate whether people rely more on the spectrum of the bouncing sounds or on the temporal pattern of the bounces.

The restitution coefficient determines a temporal phenomenon that, as we will show, plays a role in the auditory perception of the size of the bouncing ball. In the perception experiments we will focus on the ball size and the naturalness of the bouncing sounds, and keep the dimensions and material of the plate constant.

3.2 Mechanical analysis

In this section we will look at the mechanics of a bouncing ball, and analyze the mechanical variables which characterize only the plate and those which characterize only the ball parameters. First, we will try to identify the mechanical variables that depend on properties of the plate, because those are the variables kept constant and, hence, can be ignored. Using the ball parameters only we will then identify the two parameters that determine the restitution coefficient. As a last step in this analysis we will determine the influence of the plate on the spectrum of the bouncing sounds.

To investigate the relation between the mechanical variables and restitution coefficient, we must look at the time interval t_c the ball and plate are in contact with each other, because during this contact, some energy is transferred from the ball to the plate which results in v_{out} being less than v_{in} for the investigated setup. When a ball is pressed against a plate, the deformations of the ball and the plate will cause the distance of the ball center and the mid-plane of the plate to become smaller than the normal undeformed minimum distance. This change in distance, will be called α . The variation of α during a bounce is governed by a nonlinear differential equation, which was, in a somewhat different form, published by Zener [111]:

$$\frac{d^2\alpha}{dt^2} = \kappa\alpha^{1/2} \left(\frac{3}{2}\beta\frac{d\alpha}{dt} + \frac{\alpha}{m_s} \right), \quad \kappa = \frac{\sqrt{R}}{D}, \quad (3.1)$$

where κ is the contact stiffness. For each of the parameters in this equation we can identify one of two sources, the ball or the plate. Parameters for the plate are β and D , those for the ball are m_s and R . The parameter D is a plate parameter here, because the plate is much softer than the ball. The parameters of the ball, the mass and the radius, respectively, are evident; the plate parameters, the mechanical

plate impedance, β , and the elasticity coefficient, D , have been described in the previous chapter. Note that D is a plate parameter because in this setup the ball is much harder, or less elastic, than the plate. After introducing the equation and its parameters, it can be shown that only two factors determine the ball-plate contact, i.e., the contact time (for a derivation, see Chaigne and Doutaut [11])

$$t_c = 3.22 \left(\frac{m_s}{\kappa} \right)^{2/5} v_{in}^{-1/5}, \quad (3.2)$$

and the dimensionless inelasticity parameter as introduced by Zener [111]:

$$\lambda = \beta \frac{m_s}{t_c}. \quad (3.3)$$

The restitution coefficient is only a function of λ , but, due to the non-linear nature of Equation (3.1), there is, as far as we know, no analytic expression for it. The relation has been plotted graphically by Zener [111]. These formulas provide us with the insight into the relation between these mechanical parameters on the one hand, and the spectrum and restitution coefficient on the other. The restitution coefficient itself can be found by solving Equation (3.1) numerically (see, for instance, Kreyzig [53]).

The mechanical plate impedance, β , is constant when the ball is bouncing somewhere in the middle of the plate. As was shown in the previous chapter, the admittance, $1/\beta$, at the edge of the plate is about 3.5 times the admittance in the middle, which results in a lower restitution coefficient, and a different spectrum for each bounce. In our measurements of the restitution coefficient close to the edge we found a good agreement with the theory using Eichler's admittance [19]. Furthermore, we could determine that, for our plates, the admittance only changed significantly when the location where the ball bounced was closer to the edge than a few centimeters.

Once we know λ and t_c , we know the force the ball has exerted onto the plate. However, the actual spectrum of the bounces depends on the plate itself. Different plates will radiate different sounds when they are subjected to the same force. For instance, the damping in the plate has a large influence on the sound, but it has no influence on β and, therefore, neither on the restitution coefficient nor the contact time of the ball and the plate. As long as the plate is kept constant, the parameters λ and t_c determine the spectrum of the sound of a ball bouncing not too close to the border of the plate. For balls of the same material, a larger ball will bounce shorter than a smaller one.

To conclude, the size of the ball has an influence on both the restitution coefficient and the spectrum of the sound. Both restitution coefficient and the spectrum of the bouncing sound are, however, also influenced by other mechanical parameters so that listeners cannot rely on only one of them to estimate the size of a bouncing ball accurately. By combining them a listener may obtain a more robust estimate of the size of the ball but it is still impossible to calculate analytically the ball size and mass from the contact time and the inelasticity parameter without knowing other mechanical variables. We would like to know how people determine the size of a ball by listening to its bouncing sound. Do they rely more on temporal parameters, the restitution coefficient, or on the spectral properties of the sound? Indeed, do they combine those two to come to a more robust response?

3.3 General methods

In order to generate bouncing sounds, we dropped metal balls with diameters of 10, 14 or 19 mm on a plate made of MDF, which consists of wood particles that were pressed and glued together. All balls were dropped from the same height, about 10 cm above the plate. The plate had a surface of 50x120 cm and a thickness of 18 mm. The sounds were recorded using a microphone, Type Røde NT5.

An example of a recorded bouncing sound is shown in Figure 3.1b. The bouncing consists of a series of separate bounces with ever decreasing amplitude, and with intervals between bounces that decrease in length according to the restitution coefficient. It can be seen that, at least at the beginning of the bouncing series, the sound is almost completely damped out before the next bounce starts. Hence, every bounce is separated from the next by a silent period. By shortening or lengthening this silent interval, by adding or removing zero's at this position of the sound file, one can simulate the effect of a manipulated restitution coefficient. The spectral properties of each individual bounce sound change from bounce to bounce, because the reduced impact velocity results in a longer contact time according to Equation (3.2). The ball mass has, according to the same equation, also an influence on the contact time. To give an impression of the spectral consequences of the different ball sizes, the spectrum of the first bounce of each ball is depicted in Figure 3.2.

By thus increasing and decreasing the interval between the bounces, and changing the amplitude correspondingly, the restitution coefficient of the bouncing series was varied in seven steps. Step 1 corresponds to the lowest restitution co-




FIGURE 3.1: The waveforms of a bouncing ball with manipulated restitution coefficients. The ball was made of steel and had a diameter of 19 mm. The signal with the original restitution coefficient is shown in panel b. In panel a the amount of silence between the bounces and the amplitude of the consecutive bounces are lowered to simulate the effect of a lower restitution coefficient. Similarly, in panel c the silent gaps have been elongated and the amplitude increased to simulate the effect of a higher restitution coefficient. These restitution coefficients correspond to respectively step 2, 3 and 4 of the heaviest ball that was used for the perception experiment.




FIGURE 3.2: The spectra of the first bounces of bouncing balls for each of the three diameters. The spectrum of the ball with the diameter indicated above the panels is plotted in thick lines, the other two are plotted with dots for comparison.

TABLE 3.1: Restitution coefficient of the second bounce of the bouncing ball recordings. Horizontally are the five or seven manipulation steps, vertically the three different sizes.

Ball size mm	restitution coefficient step						
	e1	e2	e3	e4	e5	e6	e7
10	0.60	0.63	0.69	0.72	0.78	0.81	0.85
14	0.60	0.65	0.70	0.73	0.78	0.80	0.85
19	0.58	0.62	0.67	0.72	0.76	-	-

efficient, step 7 to the highest. The restitution coefficient varies over time, as it increases on each bounce, so the restitution coefficient of the second bounce for each ball for each manipulation step is given in Table 3.1. A picture of the waveforms with the original and modified restitution coefficients is given in Figure 3.1. By changing the radius and the mass of the ball independent from each other, the contact time and the restitution coefficient can also be changed independently, as can be seen from Equations (3.2) and (3.3). Therefore the manipulation we apply here could also result from material variations in the setup. When changing the material from lead to aluminum, and adapting the mass and radius so that the contact time remains equal, this results in a change in restitution coefficient of less than one step of our manipulation.

At the end of the bouncing process, the bounces follow each other very rapidly and there is no longer a silent gap between two bounces. To stretch the interval between the bounces, we copied a small part, from a previous bounce in the same signal, to fill the gap. In the very end, there are no longer separable impacts and this final part of the manipulated signal was simply copied from the original. Only for the largest ball it was not possible to obtain the sounds with the highest two restitution coefficients, because the silent period between the bounces was not long enough and it needed to be stretched too much, resulting in unnatural sounds. These two steps were, therefore, omitted for this ball. The restitution coefficient of the sounds changes as a function of the impact velocity, and hence over time. The recorded (original) restitution coefficient for these balls differ, and correspond to step 3 for the largest ball, step 4 for the middle ball and step 5 for the smallest ball. In this way we could manipulate the temporal pattern of the bouncing sounds, keeping their spectral properties unaffected. This resulted in $3 \times 7 - 2 = 19$ different sounds.

First, we investigated whether listeners recognized the natural relation between the restitution coefficient and the spectrum of the sound. Second, we tested if listeners do, in their judgment of the size of the ball, rely more upon the temporal information, i.e. the restitution coefficient, or upon the spectrum of the

individual sounds. The loudness of the sounds could be a very obvious cue for the size of the ball producing the sound, because larger balls make louder sounds. We are, however, interested in the spectral and temporal cues, and therefore, all stimuli were normalized for peak amplitude. Because the sounds have a very fast decay after the impact of the ball, this simple normalization effectively levels out loudness differences.

The sounds were presented in pairs to 40 participants via headphones, in a two-alternative forced-choice procedure. The sounds varied in length between about 0.7 and 2 seconds, depending on the restitution coefficient. There was 250 ms of silence between them. The levels of all bouncing sequences were normalized to have the same peak value. The largest ball produces more sound when bouncing and this would otherwise result in an obvious cue to the participants indicating the size of the ball. In Experiment 1 the participants were asked to select the most natural sound. Only pairs where both bouncing sequences originated from recording were used. Since there were seven different sounds for the two smallest balls and five for the largest ball, this resulted in $2 \times 7 \times 6 + 5 \times 4 = 104$ sound pairs. In Experiment 2 the participants were asked to select the sound coming from the largest ball. The two sounds originated from the same or different recordings. Only restitution coefficient steps 1, 3, 5, and 7 were used, resulting in 11 sounds and 110 sound pairs. In both experiments the participants could answer by clicking on one of two buttons, labeled “first” or “second”. The experiment started with six practice sound pairs for the participants to get accustomed to the sounds and the response system. The sound pairs within each of the two experiments were presented in random order. All participants took part in both experiments which took them about 20 minutes. Participants were paid 3.50 euro for participating. In the instructions the participants were told that the materials were always the same, and the ball sizes were varied. They were also told that some sounds were modified after being recorded. The participants were not explicitly told that the balls were dropped from the same height, but in principal they could derive this from the restitution coefficient and the time in between the bounces.




FIGURE 3.3: The relative preference values regarding the naturalness of different sounds are plotted as a function of the restitution coefficient step. The panels show individual data for two participants. In the left panel the data from a participant are depicted that seems to recognize the relation between restitution coefficient and spectral content of the individual bouncing sounds. In the right panel, we see the data of one participant who consistently selected the sound of longer bouncing balls as being more natural.

3.4 Experiment 1

3.4.1 Method

As mentioned, we asked the participants to choose the most natural sound out of the two of one pair. The relative preferences were calculated by averaging the number of times the participants responded to prefer a specific sound over one of the others.

3.4.2 Results

The data of each participant were analyzed in terms of the preference for each of the 19 bouncing ball sounds. Figure 3.3 shows the results for two participants. The relative preference is plotted for the 7, or 5 for the largest ball, restitution coefficient steps, and the different symbols indicate the different ball diameters. For the data of each participant a quadratic function was fitted through the relative preferences of the 7, or 5 in case of the largest ball, versions of the one sound that was manipulated to differ in restitution coefficient. For many participants these quadratic functions showed maxima located outside the range from 1 to 7 for the two smaller balls and 1 to 5 for the largest balls. For these participants, the quadratic preference function increased monotonically over the range from 1 to 7 (or 1 to 5). This means that these participants always selected the longer bouncing ball as the more natural one. In Figure 3.3b we can see one example of a partici-

(a) rc equal

s3	87	76	–
s2	83	–	18
s1	–	16	11
	s1	s2	s3

(b) rc different

s3	88	69	38
s2	79	47	19
s1	60	19	12
	s1	s2	s3

TABLE 3.2: Average percentage of “second ball is bigger” responses, for bouncing sounds with different spectral content. The spectral content of the sound is coded with “s1” meaning that each individual bounce has the spectral content of the smallest ball, through “s3” where each individual bounce has the spectral content of the largest ball. The horizontal axis represents the balls which were presented first, and the vertical axis represents the balls presented second. In the left panel the data are selected where the restitution coefficient was manipulated to be equal for both sounds. In the right panel the restitution coefficient of the first and second bouncing sound were not equal. The presentations were balanced; thus the first ball was bouncing longer as many times as the second ball. To have an equal amount of data in each cell, only the data for restitution coefficient steps 1-5 were used. The 95% two-sided confidence interval for guessing for the left panel is 44%-56% and for the right panel it is 46%-54%. Data that fall outside these intervals are printed in boldface.

pant who clearly prefers the longer bouncing ball. When, on the other hand, the maximum is located within the range of the tested restitution coefficients, the participant apparently has a preference for a particular restitution coefficient. When this is the case for all of the three ball sizes, we can say that such a participant was able to recognize the relation between the restitution coefficient and the spectral properties of the bounces. In total 12 out of the 40 participants showed such a behavior, and Figure 3.3a shows the preference data for one of them. For this particular participant the peaks of the fitted preference data are in their physically natural order, the larger the ball the lower the most preferred restitution coefficient. Furthermore, the peaks of the fitted preference curves occur, in absolute sense, close to the recorded versions. The preferred restitution coefficient steps are 2.8, 3.6, 4.2 while the natural values are 3, 4 and 5. The result shown in Figure 3.3 represent clear examples of the two different response types. Looking at the group of 40 participants we find responses that cover the whole range in between the two examples in Figure 3.3. None of the participants, however, showed a systematic preference for the low restitution coefficient steps.

(a) sizes equal					(b) sizes different				
e7	44	46	47	–	e7	46	45	41	52
e5	52	39	–	57	e5	45	46	51	54
e3	50	–	56	55	e3	49	46	49	49
e1	–	62	61	51	e1	51	48	59	50
	e1	e3	e5	e7		e1	e3	e5	e7

TABLE 3.3: Average percentage of “second ball is bigger” responses, for bouncing sounds with different restitution coefficients. The restitution coefficient was manipulated in different steps, ranging from the lowest restitution coefficient, step 1 to the highest one, step 7. Horizontally are the balls presented first, and vertically those presented last. In the left panel the data are selected for pairs in which the spectrum of each bounce was equal for both sounds. In the right panel the spectra of each bounce of the first and second bouncing sound were not equal. The order of the pairs was balanced; thus the pairs in which the first ball was the one bouncing longer were as often presented as the pairs in which the second ball was bouncing longer. To have an equal amount of data in each cell, only the data from the smaller two balls were used. The 95% two-sided confidence interval for guessing for the both panels is 43%-58%. For the left panel, results do not differ significantly from guessing ($\chi^2_{df=11} = 9.02$; $p = 0.62$), or a 50% score in each cell. For the right panel, the results do not differ significantly from guessing ($\chi^2_{df=15} = 4.64$; $p = 0.995$).

3.5 Experiment 2

3.5.1 Method

In experiment 2 the participants were asked to judge the size of the bouncing ball. We again used a two-interval forced-choice procedure and the participants had to indicate which of the two sounds was produced by the larger ball. In contrast to experiment 1, for this test, the two presented sounds were recordings from different balls, and the restitution coefficient was varied for each of them. The results are shown in Table 3.2 and 3.3. In Table 3.2a, we can see that people on average label the sounds with the spectral content of a larger ball as the larger ball. In Table 3.2b we can see that average performance does not degrade, even if the restitution coefficient differs between the two sound samples. On the diagonal of this same table, we can see another effect. If both balls are equally big, we would expect an equal division between choices for the first and second sample. As we can see on the diagonal of Table 3.2b this is not the case. When both balls are small, the participants more often labeled the second ball as larger. When both balls are large, the participants more often labeled the first ball as larger. Al-




FIGURE 3.4: Average response to the question which of the two presented sounds of a pair represented the larger ball. The data are analyzed for each participant in terms of differences in spectral content of the bounces. A “first” response is coded with -1 and a “second” response is coded with 1. On the unitless y-axis are the averages of these coded responses when both sounds differed in spectral content of the individual bounces. The x-marks represent the data where the first sound had the spectral content of the larger ball, the circles represent the data where the second sound had the spectral content of the larger ball. The order in which the sounds differing in restitution coefficients were presented was balanced, i.e., pairs in which the first ball was bouncing longer were as often presented as the pairs in which the second ball was bouncing longer. If all responses were based upon the spectral information, the x-marks should be at -1 and the circles at +1.

though we have not proven this hypothesis, the following reasoning may explain such behavior. When the listener concludes that the sound of the first ball, he or she already estimates its size. If he or she hears the first ball is the smallest, he or she hypothesizes that the next ball will be larger and responds this way unless he or she is proven wrong by the second sound.

The data from Table 3.2 were recalculated for Table 3.3, but now analyzed as a function of the restitution coefficient of each of the two balls in a pair. Again, the table shows the average response to the question which of the two presented sounds of a pair represents the larger ball. We can see most entries are close to 50% indicating that the participants, on average, do not respond very much to differences in the restitution coefficient. In contrast, individual participants do respond to the restitution coefficient but there seem to be large differences between participants as we will see in the next paragraph.




FIGURE 3.5: Average response to the question which of the two presented sounds of a pair represented the larger ball. The data are analyzed for each participant in terms of differences in restitution coefficient of the two bouncing balls. A “first” response is coded with -1 and a “second” response is coded with 1. On the unitless y-axis are the averages of these coded responses when both sounds differed in restitution coefficient. The x-marks represent the data where the first sound was the longer bouncing, which is usually the smaller ball. The circles represent the data where the second sound is the longer bouncing. The order in which the sounds differing in spectral content were presented was balanced, i.e., pairs in which the first ball had the spectral content of the larger ball were as often presented as the pairs in which the second ball had the spectrum of the larger ball. If all responses were based upon the temporal information, the x-marks should be at -1 and the circles at +1.

In Figures 3.4 and 3.5 the individual performance of each participant is depicted. The responses labeled “first” and “second” are coded -1 and 1, respectively. In Figure 3.4 and 3.5 the averages of these responses are shown for different subsets of the data. In Figure 3.4, the data where the first sound has the spectrum of a larger ball is plotted with a cross and the data where the second sound has the spectrum of the larger ball is plotted with a circle. All different combinations of the restitution coefficient were used, i.e., the first is either bouncing shorter or longer than the second or they could be bouncing equally long. In order to limit the number of stimuli, only four steps for the restitution coefficient were used, steps 1, 3, 5 and 7. The used data are counterbalanced for the restitution coefficient, thus the first sound is as often from the longer bouncing ball as the second. This means that, if the line between these two points is long, the participant has used the spectral information in the sound to come to the answer. For Figure 3.5, the data where the first sound had the lower restitution coefficient of a larger ball are plotted with a cross and the data where the second sound had the lower restitution coefficient is plotted with a circle. All different combinations of the spectral information were used, and the used data are counterbalanced for the spectral information. This means that, if the line between these two points is long, the participant has used the restitution coefficient in the sound to come to his answer. The two plots are not based on the same data, since the data where both balls have the same spectral content were not used in Figure 3.4 but they were in Figure 3.5. The opposite is true for sounds with the same restitution coefficient; data where both balls have the same restitution coefficient, were not used in Figure 3.5 but they were in Figure 3.4. For both plots, when the line is not centered around zero, this means that there was a response bias, in the sense that the listeners preferred the either the first sound or the second sound more often.

From Figures 3.4 and 3.5 we can draw the following conclusions regarding the different ways the participants responded to the stimuli. Participant number 1 has a very limited difference in response depending on the temporal information, the line connecting x and o is very short in Figure 3.5, but spectral information has a large effect on his or her answers. We can conclude that this participant relies on spectral information to come to his or her answer. For participant number 3, however, the situation is reversed. He or she almost does not respond to spectral differences, but very much to temporal differences. Yet another situation we see for participant number 15 who seems to base the judgment neither on temporal nor on spectral differences. There is one inverse responder, participant number 6, for the spectral cues, and there are 12 inverse responders for the restitution

coefficient, among whom number 6. These people do not recognize the natural relation between the spectral content or the restitution coefficient and the ball size.

3.6 Discussion and conclusion

We have analyzed the differential equation governing the restitution coefficient and found a relation in the contact force between the spectral and temporal properties. In a perception experiment, we found two types of responses, some participants were able to identify some relation between the spectral and temporal properties, and other participants chose always the highest restitution coefficient as most natural. This is similar to the division in two groups was found by Canévet *et al.* [10] when asking subjects for the pleasantness of single impact sounds. In our case, there is no clear distinction between two groups, most people respond somewhat in between the two described types of responses.

Second, we argued that both the spectral and temporal cues are also influenced by other parameters such as the plate thickness. When asked for the size of the ball, the participants respond, on average, according to the spectral content of the sound. The restitution coefficient had an influence on the responses for many participants, but in a less consistent manner. Such a listening strategy can be successful in an experimental setting like the one we used, where the plate parameters are kept constant. It should, however, lead to wrong judgments when other mechanical parameters, like the plate thickness, are variable. In such a situation, only a combined evaluation of spectral and temporal cues would allow to correctly judge ball parameters, because each one individually is affected by variations in both ball and plate parameters.

4 Temporal Aspects of Rolling Sounds: A Smooth Ball Approaching the Edge of a Plate¹

Abstract

We measured the sounds of smooth metal balls rolling over medium density fiberboard (MDF) plates. In the spectrograms of these sounds we observed gradually varying ripples. These ripples were more closely spaced for the sound generated in the middle of the plate than for the sound generated closer to the edge. Furthermore, the spacing for lower frequencies was somewhat closer than for higher frequencies. It is shown that this pattern arises from the interference between the sound directly generated at the point of contact between ball and plate, and the sound reflected at the edge of the plate. This effect was added to synthesized rolling sounds which resulted in a more natural sound. A discussion is presented concerning the perceptual relevance of this pattern.

4.1 Introduction

This chapter is concerned with a spectro-temporal pattern in the spectrogram of the sound of a ball rolling over a plate. This study is done in the wider context of a study into the human capabilities of extracting physical and geometric object properties, such as size and shape, from the sounds these objects generate when impacted by other objects. Various authors (e.g., [22][56][54]) have shown these capabilities. Until now, they used metal objects with simple geometric forms, and the sound was induced by simple impacts. The acoustical analysis of these systems, as given in these papers, is relatively straightforward; the impact sounds are modelled as a sum of modal frequencies with exponentially decaying en-

¹This chapter is based on Stoelinga, Hermes, Hirschberg and Houtsma [98]

velopes. This was further developed into algorithms to synthesize more complicated impact sounds such as the sounds of balls or pebbles rolling in metal vessels and other hollow metal objects (Van den Doel *et al.* [105]). The results were informally compared with recorded sounds of such systems.

The system we studied consisted of a ball rolling over a wooden plate. Due to strong damping and the absence of clear modal frequencies, the method by Van den Doel *et al.* could not be applied. Houben *et al.* [41] demonstrated for a wooden ball rolling over such a wooden plate, that human listeners are capable of distinguishing small from large and slowly from rapidly rolling balls on the basis of their rolling sounds. By signal analysis techniques they then manipulated the spectral and temporal properties of these sounds and could derive from perception experiments with these manipulated sounds that both temporal and spectral properties played a role in the successful judgment made by the listener. They additionally showed that quasi-periodic amplitude modulation, produced by the not perfectly spherical shape of the wooden balls, when audible, played a dominant role in the judgments by the listeners. For the rest, Houben *et al.* could not directly link the spectral and temporal signal properties on which the listeners based their judgments with the physical properties of the system that determined these spectral and temporal properties. An obstacle to this was the lack of an acoustic model of the rolling-ball system.

We believe that an appropriate understanding of the acoustic properties of the system is necessary, if we want to understand the temporal and the spectral information in the sound signal that the human listeners use in making judgments about the physical properties of the system, such as the size, material, or velocity of the ball rolling over the plate.

Other studies on rolling sounds (Thompson *et al.*[101][100]; Scheuren[88]) focused on the road and railway noise produced by cars and trains. The primary interest of these studies was noise reduction and, hence, they mainly focused on the distribution of acoustic energy over the spectrum. Hardly any attention was paid to temporal and more complicated properties of the noise. Moreover, the physical systems are essentially different from a system consisting of a ball rolling over a plate. For instance, it will be shown that the presence of edges is very important in modeling the acoustics of the ball-plate system.

This study is focused on a spectro-temporal pattern in the spectrogram of the rolling sounds of a smooth ball approaching the edge of a wooden plate. The origin of this pattern will be shown to lie in the reflections at the edge of the plate. This effect is first studied by means of cross-correlation techniques, first on un-

processed accelerometer signals, and then on band-pass filtered versions of these signals. This is used to determine the wave velocity for a range of frequencies. This is of significance since the waves predominantly consist of bending waves which obey the fourth order dispersive wave equation. For these waves the propagation velocity is frequency dependent, which will appear to be important and this effect will be considered in more detail. A simulation model will be proposed, which is based on signal analyses of recorded sounds. In this model, the mentioned spectro-temporal effects of a ball rolling towards the edge of a plate are included on the basis of a simplified physical analysis. The perceptual relevance of this pattern will be presented.

4.2 Measurement setup

The vibrations of the plate induced by the rolling ball were recorded using the setup shown in Figure 4.1. The balls were made of stainless steel and had diameters of either 35 or 55 mm. The plates, 49 cm wide and 122 cm long, consisted of MDF, wood fibers compressed and glued together. Their thickness was 6, 12 or 18 mm. The results used as illustration in the figures of this paper were obtained with the thickest plate.

The setup rested on a table from which it was isolated by a 25 mm thick layer of soft foam. This insulation layer was necessary to avoid vibrations of the table. The effect of the plate's support on our results has been checked by carrying out some measurements on a plate supported by four small soft balloons filled with air. While a reduction in damping of low frequencies is observed the relevant spectro-temporal effects discussed further were not altered. This indicates that the 18 mm thick plate can be considered as "free". Especially for the 6 mm plate, however, the extra damping caused by the foam was very prominent. A possible consequence of this extra damping is a complex wave velocity [30, p. 173].

In order to give the balls a well defined velocity along the center line of the plate, they were rolled from a slide onto the plate. The height of the slide was 25 cm. The balls were released from various heights on the slide to vary their velocity. In order to measure the velocity of the ball, its course interrupted the beams of six independent light-gates, placed at intervals of 20 cm. The measurements showed that, during one run, the ball's velocity was constant within a few percent. The range of velocities we used in the experiment was varied from 0.6 to 1.5 m/s.




FIGURE 4.1: Side-view of the experimental setup: the plate supported by a foam layer is placed on a table. The slide is separated from the plate by a narrow gap. Light-gates are used to measure the velocity of the ball.

The slide was separated from the plate by a narrow slit, about 0.5 mm in width, to avoid the transmission of vibrations. The slide had a smooth bend near its end to transform the vertical velocity of the balls into a horizontal one. At the other side of the plate, the ball was left free to roll off.

This experimental setup was chosen in order to prevent bouncing or amplitude modulation by less perfectly round balls. Furthermore, relatively heavy, polished metal balls were used, which resulted in a smooth and stationary rolling sound. In this way, the spectro-temporal pattern could best express itself and was not obscured by bouncing ticks or other amplitude fluctuations. In fact, the sound generated in this manner does not sound very much like rolling. When the ball is gently slid over the plate, while preventing it from rolling, a similar sound is generated. Hence, the analysis presented here is equally valid for balls sliding softly over the plate.

4.3 Acoustical analysis

An important variable in the rolling-ball system is the surface roughness of the plate. Surface irregularities of the ball will be ignored, since the ball was made of stainless steel and polished. As long as the velocity of the ball is low, the ball will continuously keep contact with the plate. Due to the surface irregularities it will then move up and down and the corresponding reaction force will excite the plate continuously. When the velocity of the ball gets higher, it may occasionally lose contact with the plate, causing light impacts when it comes down again. Very




FIGURE 4.2: Average spectrum of a ball rolled over the plate.

rough irregularities may also induce such bouncing-like rolling behavior at lower speeds. The smooth plates we used did not cause such a bouncing.

Even though the ball is in continuous contact with the plate, the interaction is of a non-linear nature. The implications of these non-linear effects have not been fully investigated, but it is assumed that the ball induces a force with band-pass characteristics onto the plate. The plate surface will not contain very low frequencies, since it was polished. Spatial irregularities that are very close together do not move the ball very much either, because it simply rolls over them. Due to these two effects the force that the ball exerts on the plate will have a band-pass character. It can easily be seen that the forcing frequency Ω depends on the velocity of the ball V for one wavelength λ of the roughness spectrum,

$$\Omega = \frac{2\pi V}{\lambda}. \quad (4.1)$$

Therefore the frequency range of this band-pass filter scales with the velocity of the rolling ball.

Apparently, the spatial-frequency characteristics of the plate roughness and the velocity of the ball determine the bandwidth of the frequencies by which the plate will be excited. In the system we used, this bandwidth was rather broad. An example of an estimated spectrum is shown in Figure 4.2. It was calculated by averaging the spectra of several successive windowed segments of the signal. This averaging was done to reduce the variance of the spectrum. This method results in what is known as Welch averaged periodograms [79].

From the absence of clear peaks one can see that no strong modal frequencies were found. This is due to the strong damping inside the plate. When waves are

induced at one end of the plate they have almost vanished when they reach the other end of the plate. On the other hand, reflections can be found when the ball is near the edge of the plate, which will play an important role in the remainder of this chapter. (see section 4.4).

Another important aspect relates to the kind of waves playing a role in the rolling ball system. When a plate is hit by another object, vibrational waves spread from the point of contact over the plate. Following classical plate theory, the vibrational waves are bending waves which are described by a fourth order differential equation. In this case the group velocity v_g is not independent of frequency but is proportional to the square root of the frequency [48]:

$$v_g = 2 \left(\frac{D}{\rho h} \right)^{1/4} \omega^{1/2}, \quad (4.2)$$

where D is the flexural rigidity of the plate, h the plate thickness, ρ the plate density and ω the angular frequency.

4.4 Reflections

In a string or membrane the theoretical description of reflections at an open end is relatively straightforward. Some analytically calculated examples of the reflections for bending waves in bars and for waves reflecting at the surface of a semi-infinite medium indicate, however, that the analysis of the reflections at the free end of a plate is more complicated than that of a membrane (see [30, section 4.2]). Instead of deriving the behavior of the reflections mathematically, their properties will be determined by calculating the correlation function from measured data, in a similar way as treated by Bendat and Piersol [5]. In two-dimensional structures with reflecting edges there are many different paths consisting of straight lines running from the excursion point to the measuring point: direct and via one or more edges. When trying to measure these reflections with the auto-correlation function, $R_{ff}(\tau)$, the peak at $\tau = 0$ was found to be very strong while others were much weaker and difficult to detect precisely. This could be attributed to the dispersion and damping of the waves as they travel across the plate. This was partly solved by using two measuring points instead of one and determining the cross-correlation function between the signals measured at these points (see Figure 4.3). This correlation procedure was repeated for various bandpass-filtered versions of the signals, which also increased the sharpness of the figures. For two




FIGURE 4.3: Traveling paths from the ball (marked with a solid circle) to the accelerometer (marked with a square) with a single boundary, resulting in a mirror image of the ball (marked with a dashed circle) and four different paths from the source and its mirror to the accelerometers.

accelerometer signals $f(t)$ and $g(t)$, the cross-correlation function is defined as

$$R_{fg}(\tau) = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T f(t)g(t - \tau)dt. \quad (4.3)$$

Consider a semi-infinite plate with a single edge as shown in Figure 4.3. Positioned on the plate are the two accelerometers and a ball impacting the plate. For each accelerometer there are two paths running from the ball to the accelerometers: one direct path from the excitation point; the other reaching the accelerometer after being reflected at the edge. This second path can be modeled as coming from a virtual mirror image at the other side of the edge. The direct wave $f_0(t)$ arrives first, followed by the reflection from the mirror image. Due to the spreading of the wavefront over a larger region there is a reduction in amplitude given by α_1 and α_2 ,

$$f(t) = \alpha_1 f_0(t - \tau_1) + \alpha_2 f_0(t - \tau_2), \quad (4.4a)$$

$$g(t) = \alpha_3 f_0(t - \tau_3) + \alpha_4 f_0(t - \tau_4). \quad (4.4b)$$

Using Equation (4.3), this cross-correlation function splits up into the weighted sum of four of the original auto-correlation function shifted in time. These auto-correlation functions have peaks at $\tau_3 - \tau_1$, $\tau_3 - \tau_2$, $\tau_4 - \tau_1$ and $\tau_4 - \tau_2$, respectively. If these peaks are spaced widely enough, they can be distinguished separately in the cross-correlation function. As more edges are added, the number of mirror images increases accordingly. For an actual plate with four edges, four mirror images will be found. But these are only mirror images of the first order, repre-

senting waves reflected at one edge. Similarly, waves reflected more than once can be considered as coming from mirror images of higher order, resulting in a cross-correlation function consisting of the weighted sum of a large number of time-shifted auto-correlation functions. As in the MDF plates we used the damping was quite high, the weight factors $\alpha_i \alpha_j$ decreased rapidly with increasing path length. Hence, we will only consider the first-order reflections.

So far, only a fixed excitation point was considered. To obtain useful information about the rolling ball this must be extended to a moving excitation point. In order to study the signal at time t_0 , a short segment is gated out by multiplying the signal by a window function $w(t)$ shifted t_0 in time. In each windowed segment, the excitation point, or ball position, is considered fixed. For each t_0 , we then calculate the cross-correlation function of this windowed segment, or, expressed mathematically,

$$R_{fg}(\tau, t_0) = \frac{1}{E_w} \int_{-\infty}^{\infty} f(t)g(t - \tau)w(t - t_0)^2 dt. \quad (4.5)$$

A non-overlapping square window function was used with a length of 4.167 ms. In the range of the measured velocities, this corresponds to rolling over a distance of between 2.5 and 6.25 mm. E_w represents the energy of the signal within this window.

It appeared that during rolling the energy of the signal was very unevenly distributed over time, so that one windowed segment could contain a lot more energy than the next. In order to compensate for this, the cross-correlation functions were normalized for energy. We will indicate the so obtained figures with *running cross-correlogram*.

The two-dimensional function given by Equation (4.5) was then plotted with t_0 on the horizontal axis, and τ on the vertical axis. Hence, for each t_0 considered, the cross-correlation functions were displayed vertically. Their values were converted into a gray scale with higher values being lighter than lower values. An example is presented in the right panel of Figure 4.4

When the ball moves over a certain distance, the mirror images move along over the same distance. Depending on the configuration, this can lead to changing differences in path lengths between the excitation point and its images, and the accelerometers. This will result in changes in the position of the peaks in the cross-correlation functions. On the other hand, the position of the ball, its mirror images and the accelerometers are known. Hence, the actual differences in path lengths can be measured. And if, for the moment, it is assumed that




FIGURE 4.4: The pattern in the running cross-correlogram, as calculated (left) and measured (right). The two accelerometers were mounted at the edge of the plate on $1/3$ and $2/3$ of its length. The calculated lines approximate the measured ones.

the wave velocity is constant, the time differences with which the waves travel these distances can be compared with the positions of the peaks in the measured cross-correlation functions. This provides us with a way to verify whether or not, in the complicated situation of a real plate with its dispersion and its damping, these reflections are very different from the simplified theoretical cases.

In the left panel of Figure 4.4 we plotted the calculated maxima of the running cross-correlogram under the assumption that the waves travel with a velocity of 300 m/s. These lines are only plotted if the traveling distance from the origin, and hence the attenuation of the signal, is less than some specified value. Energy loss at the reflections was not taken into account.

The right panel displays the actually measured running cross-correlogram. If the position of the lines in the upper panel is compared with the light bands representing the maxima in this running cross-correlogram, the correspondence is indeed remarkable.

Note that in this simplified description it is assumed that the waveform is not distorted too much in the course of traveling through the plate. This implies that the wave velocity is more or less constant, and that issues like dispersion and energy loss at the edges have not yet expressed themselves in clear changes in shape of the wavefront. The correspondence between the two patterns described in the lower and upper panel of Figure 4.4 shows that, to a certain extent, this is indeed fulfilled. This can partially be explained by the limited frequency range of the excitation. So it can be concluded that the reflected waves are strongly correlated with the original waves. It will be shown later that the resulting similarities

between original and reflected waves are the origin of the interference pattern as found in the spectrogram of the rolling sounds.

4.5 Wave velocity

In the previous sections the wave velocity was regarded as constant. To calculate maxima in the patterns seen in the running cross-correlogram, we had to estimate the wave velocity. Now things can be turned around. One may try to estimate the wave velocity from the running cross-correlogram. This can be done by calculating, for a plausible range of velocities, the correspondence of the calculated lines of maxima with the actual maxima in the measured running cross-correlogram. This has been done informally by printing the calculated lines on top of the measured running cross-correlograms, or by summing, for the range of velocities considered, the values of the running cross-correlogram underneath the calculated lines. The highest value is then expected to represent the best estimate of the wave velocity. This procedure still assumes that the wave velocity is constant within the frequency band of the acoustic signal, and it was shown that possible discrepancies did not distort the reflected signal too much. Classical plate theory predicts, however, that the traveling-wave velocity is proportional to $\sqrt{\omega}$.

Dispersion of the waves leads to a change in shape of the reflected wave, and thus the correlation technique will have to be adapted. In practice the correlation function is blurred due to this dispersion, which means that instead of well defined narrow and high peaks, smooth broader and lower ones will be found. The amount of blurring depends on the range of frequencies in the waves and the distance they have traveled.

In order to get a frequency dependent estimate of the wave velocity, we will now carry out the above mentioned procedure for a series of band-pass filtered signals. The centre frequencies of the filters ranged from 1 to 10 kHz. Indeed, Figure 4.5 shows two running cross-correlograms, one for a center frequency of 774 Hz and one for 5995 Hz. As mentioned, we then calculated, for a range of velocities, the fit of the maxima predicted for these velocities with the actual maxima in the running cross-correlograms. For low frequencies, the wavelength is no longer small compared to the travelling distances, and this method loses its accuracy. The measured values are plotted in Figure 4.6 from 1 kHz (wavelength about 25 cm) to 10 kHz.




FIGURE 4.5: The running cross-correlograms of signals that were first filtered with a bandpass filter. The centre frequency of the filter was 774 Hz for the upper panel and 5996 Hz for the lower panel. The pattern is compressed in the y-direction in the case of the highest band-filter, due to the higher wave velocity for higher frequencies.




FIGURE 4.6: The bending-wave velocity can be estimated by matching measured data with analytically calculated lines of the filtered running cross-correlogram. These estimated values are indicated here with a circle. The line indicates the theoretical group velocity obtained using the theoretical wave velocity, from the statically measured Young's modulus.

To verify the plausibility of these values, the Young's modulus and density of the MDF were measured. The Young's modulus was measured by statically loading a strip cut from the plate with various weights, in the same direction as the plate is distorted by bending waves. This is the direction in which the wood particles are pressed together, and Young's modulus measured in this way (5.7 kN/mm^2) is somewhat higher than the (true) Young's modulus, measured in other directions (4.2 kN/mm^2). From these measurements the theoretical wave velocity can be calculated. The phase velocity is also measured by determining the frequency of the first resonance mode of a smaller plate. This smaller plate, as well as the strips used for the measurements of the Young's modulus were sawn off the plates used to generate the rolling sounds before the experiments. The resulting wave velocity at the first mode of this smaller plate (52 Hz) is 50 m/s, quite close to the theoretical value based on the statically measured Young's modulus (56 m/s).

In Figure 4.6 the group velocity, as calculated via the statically measured Young's modulus is compared with those measured via the correlation method, described above. As can be seen, there appears to be a systematic error increasing with frequency, resulting in measured velocities that are about half as high as the theoretical values based on the statically measured Young's modules. This could be caused by, for instance, internal friction or dispersion, two aspects that are not included in the model.

The velocity found when we do not filter the signal before determining the cross correlate (see Figure 4.4) is indeed about equal to the velocity of strongest frequencies (see Figures 4.2 and 4.6).

4.6 Comb-filter model for reflections at the edge

In the previous section it was shown that, in our configuration of a metal ball on an MDF plate, the waves, generated at the point of contact between ball and plate, travel through the plate obeying classical plate theory. Moreover, these waves reflect at the edges of the plate and interfere with the direct, unreflected waves. Damping and dispersion are, on the one hand, so high that reflections of higher order do not play a role of significance. On the other hand, the distortion is not so strong that the similarity between direct and reflected waves is lost.

The presence of reflections can be described on the basis of virtual sound sources, so-called mirror images as already depicted and described (see Figure 4.3). Let us for the moment consider only two mirror images moving in direc-




FIGURE 4.7: The position of the two most important mirror images considered in the comb-filter model. Indicated are the directions in which the ball and its images rolls, as well as the distances, l and L , used in the formulas.




FIGURE 4.8: Model for the effect of the images. The two delayed and damped versions of the direct waves are added to this direct wave, due to the effect of reflecting at the boundaries of the plate.

tions opposite to the original ball as shown in Figure 4.7. One mirror image is positioned at the other side of the edge the real ball leaves behind, while the second mirror image approaches the real ball from the other side of the edge that the ball rolls to. This can be modelled as shown in Figure 4.8. The delay between direct waves and reflected waves can be calculated by:

$$\tau_1(t) = l/c = (vt)/c, \quad (4.6a)$$

$$\tau_2(t) = (L - vt)/c, \quad (4.6b)$$

where c is the group velocity of the vibrational waves in the plate, v the velocity of the ball, l is the distance from the mirror image to the real plate, and L is the length of the real plate. The attenuation factors α are also time dependent. They decrease for mirror images moving away from the real ball and increase for mirror images approaching the ball. If dispersion and damping are neglected, they can be approximated by

$$\alpha_1(t) = \alpha_0 l = \alpha_0 vt, \quad (4.7a)$$

$$\alpha_2(t) = \alpha_0 (L - vt). \quad (4.7b)$$

In the configuration we used, it appeared that the damping was so large that the α 's almost vanished in the middle of the plane. As a consequence, the model can be simplified by only considering one boundary at a time.

Let us now look at the consequences of the reflections on the spectrum of the rolling sound. We start by considering a fixed position of the source, and hereby we eliminate the time dependence of τ_i and α_i . The complex spectrum of the direct sound will be indicated with $X(\omega)$; the direct sound plus the reflected sound will be indicated with $Y(\omega)$. If just one reflection is considered, as argued in the previous paragraph,

$$Y(\omega) = X(\omega) + \alpha_i e^{-j\omega\tau_i} X(\omega), \quad (4.8)$$

the frequency response function $H(\omega)$ and the power spectrum $|H(\omega)|^2$ can be calculated as

$$|H(\omega)|^2 = 1 + \alpha_i^2 + 2\alpha_i \cos(\omega\tau_i), \quad (4.9)$$

with $i = 1$ or 2 for the desired edge.

This shows that the power spectrum of the combined sound can be calculated from the spectrum of the direct sound by multiplying the spectrum of the direct sound by a constant plus a cosine. This results in a rippled spectrum. Due to the fact that the vibrational waves travel at a much higher speed than the ball, we can still consider the position of the source constant, and use the given $H(\omega)$, while calculating the temporal variation of this rippled spectrum. Using Equations (4.6) and (4.9), it can be shown that the minima of this ripple are located at:

$$\omega = \frac{(2n + 1)\pi c}{l}, \quad (4.10a)$$

$$\omega = \frac{(2n + 1)\pi c}{L - l}, \quad (4.10b)$$

with $n = 0, 1, \dots, \infty$.

When we do not, for the moment, consider the frequency dependence of the wave velocity, the power spectrum for one fixed τ is plotted as the dashed line in Figure 4.9. As a consequence of this frequency dependence of the wave velocity, $c = \sqrt{\omega}c'$, the time delay τ_i in between the arrival of the direct and reflected wave is not a constant value but frequency dependent

$$\tau_i = \frac{\tau'_i}{\sqrt{\omega}}, \quad (4.11)$$

resulting in the power spectrum

$$|H(\omega)|^2 = 1 + \alpha_i^2 + 2\alpha \cos(\sqrt{\omega}\tau'_i), \quad (4.12)$$

and now the minima of this ripple are located at:

$$\omega = \left(\frac{(2n + 1)\pi c'}{l} \right)^2, \quad (4.13a)$$

$$\omega = \left(\frac{(2n + 1)\pi c'}{L - l} \right)^2. \quad (4.13b)$$

The pattern of the minima in the power spectrum can be scaled in two ways. First, increasing the size L of the plate results in a wider pattern. Second, varying c/l , for instance, by varying the thickness of the plate, changes the steepness of the patterns. Note that the patterns do not depend on the ball velocity but on the ball position.




FIGURE 4.9: Power spectrum of the model for the effect of the images, here drawn for $\alpha = 1$ and a fixed wave velocity (dashed line) or a frequency dependent wave velocity (solid line). The minima depend on τ , and τ itself depends upon the ball position.




FIGURE 4.10: In the spectrogram of a ball rolling from one side to another over an approximately free plate, U-shaped strokes of a lighter color appear, indicating less power at these frequencies. These are drawn schematically in the top panel and calculated from measurements in the bottom panel.

This final result is compared with measurements in Figure 4.10, where the minima are depicted with lines in the upper panel and are seen as strokes of a lighter color in the lower panel. Due to the earlier mentioned effects as dispersion and damping, the waves are not completely canceled at the minimum, which would result in completely white strokes. Note that τ is considered as time independent within one analysis window of the spectrogram, and changes between different time windows. The same restraint is valid for the synthesis algorithm discussed in the next section.

4.7 Synthesis

Next, we tried to get an impression of the perceptual effect of this rippled spectrum, by simulating the effect on synthesized noise and comparing it with the original noise. The synthesis algorithm was divided in two steps. We have started from Gaussian white noise and filtered it in such a way that its long-term power spectrum matched the average power spectrum of recorded rolling sound. By doing so, it is assured that all the spectral properties due to ball and plate properties as well as the velocity of the ball, are well represented in the obtained sound. Also spectral effects which are constant in the course of rolling, for instance due to the radiation from the plate, are taken into account by this step. All temporal properties, however, are ignored by this procedure. Hence, the direct signal consists of band-pass filtered noise. The noise coming from the mirror images can then be added to the original noise according to the attenuated delay modeled in Figure 4.8. This was first simulated for a frequency independent wave velocity. In this case the contribution of the reflections to the resulting sound was unnaturally strong overdone. This is a consequence of a very prominent so-called repetition pitch in this simulation. The perceptual background of repetition pitch will be described below. Adding more mirror images in the simulation did not result in any further improvement in the naturalness of the simulated rolling noise. A considerable improvement was obtained, however, when the simulation included that the wave velocity is frequency dependent. The perceptual effect of the too prominent repetition pitch disappears. The effect of the spectro-temporal pattern is still well audible, and results in a considerable improvement of the naturalness of the sound.

4.8 Repetition pitch of echoed noise

The study of repetition pitch has been of great interest to pitch-perception research. Repetition pitch can be perceived when a signal, generally a noise signal, is delayed and added to itself. The frequency of the pitch perceived corresponds to the inverse of the time delay between the two signals. Yost and Hill [108] used a model, similar to the one shown in Figure 4.8, to generate these repeated-noise signals. They concluded that a pitch can be detected even if one signal is up to 20 dB weaker than the other, or, applied to our situation, the reflection can be up to 20 dB weaker than the direct signal. Some signal alterations have been applied to the delayed signal such as filtering (Yost [109]) or phase shift (Bilsen [6]), and this is shown not to influence the presence of a repetition pitch. However, the perceived prominence of the pitch may decrease with these alterations of the signal.

In our situation of the ball rolling over the plate there are some differences with the repeated noise used above. First, the delay between the two noise sources varies with time. If the wave velocity would be independent of frequency, this would result in a time-varying repetition pitch. When the frequency dependence of the wave velocity is taken into account, the situation is more complex. In that case, the ripple on the power spectrum is no longer strictly periodic (see continuous line in Figure 4.9). Hence, a well defined repetition pitch can no longer be perceived. Rather, the non-periodic ripple colors the spectrum in a more indefinite way. When the ball approaches the edge which reflects the waves, the ripples on the spectrum become more widely spaced. Subjectively, one can hear that something gets higher. But since the ripple on the spectrum is not periodic, this, now, is not perceived as a strict increase in pitch. Neither is it perceived as a change in brightness or sharpness, as these measures are defined as the balance between high and low frequencies, which remains constant in our case (for a definition of the various sound qualities see [70][112]). Furthermore this effect becomes stronger when the ball comes closer to the edge, because the reflections are strongest there.

Note that the interference pattern consists of a change in time of a non-periodic ripple over the spectrum. Integration, either in the time or in the spectral domain results in obscuring or vanishing of this interference pattern. Hence, if human listeners indeed extract information from this interference pattern as to the position and the velocity of the rolling ball, this information resides essentially in the spectro-temporal domain.

4.9 Discussion

In our study of the sound of a metal ball rolling over a wooden plate, we encountered a spectro-temporal pattern consisting of a time-varying spectral ripple. The spacing of the ripples was not equidistant and decreased with distance from the edge of the plate. By predicting and measuring the running cross-correlogram between two accelerometer signals, it was shown that this pattern was due to the interference between the direct waves generated at the point of contact between ball and plate, and the waves reflected at the edge of the plate. Due to the strong damping in the plate, only first-order reflections had to be taken into account. Higher-order reflections only played a minor role. In addition, by calculating the running cross-correlograms for different frequency bands, it was shown that the travelling wave velocity was not independent of frequency. In accordance with traditional plate theory the traveling wave velocity increases with the square root of frequency. On the basis of this result, the spectro-temporal interference pattern could quantitatively be described.

It was argued on the basis of the results of pitch-perception research that the human listener is very sensitive to the kind of information present in this spectro-temporal interference pattern. The acoustic vibrations induced by a ball rolling over a wooden plate contain a spectro-temporal pattern listeners can use in estimating the time needed for the ball to reach an edge of the plate. Research by Cabe and Pittenger [9] and Lee and Reddish [59] showed that such information can indeed be gathered from similar properties of the visual or auditory signal. Using this information together with other sources of acoustic information, such as the spectral centroid [41], the spectral tilt, and the temporal variations correlated with the angular speed of the ball (Houben and Stoelinga [44]), allows the listener not only to estimate physical properties of the ball such as its velocity and its size, but also plate properties such as its material and its dimensions, which can never be perceived on the basis of one parameter alone.

In a following set of perception experiments we will investigate if the information conveyed by this spectro-temporal interference pattern is actually used by human listeners to determine properties such as the ball velocity or the time to reach the edge of the plate. We will do this on the basis of sounds synthesized by simulating the process of direct sound reflected at the edges of a plate. By varying this pattern independent of other physical parameters, which in the natural situation are necessarily coupled and thus correlated with this pattern, we will determine whether indeed the information present in this pattern is used

by human listeners in reconstructing an auditory image of what happens around them.

In summary, if accelerometers placed on a wooden plate record the vibrations induced by a metal ball rolling over this plate, the running cross-correlogram of the recorded signals indicates the occurrence of reflections of the waves at the edges of the plate. These running cross-correlograms additionally allow the measurement of the frequency dependent wave velocity in the plate. Moreover, the interference between the direct waves coming from the point of contact between the ball and plate and the waves reflected at the boundary cause a frequency-dependent ripple over the spectrum of the signal. Since the spacing of the ripple becomes wider as the ball approaches the boundary of the plate, the spectrogram shows a time-varying interference pattern. This effect can be described quantitatively by modeling the reflections as coming from mirror images outside the plate. This interference phenomenon is expected to be a source of information the listener uses to estimate the position and the velocity of the rolling ball. Including this interference effect in the synthesis of rolling-ball sounds on the basis of such a simple model indeed improves the perceived naturalness of the sound. Especially the frequency-dependence of the wave velocity in the model contributes to this naturalness. As this frequency dependent velocity depends on plate properties such as its material and thickness, this interference pattern may additionally provide the listener with perceptual information for these plate properties. Hence, such models provide the basis for further perception research on rolling sounds.

5 | Influence of Wave Reflections at the Plate Edge on Perceiving the Rolling Direction of a Rolling Ball

Abstract

When a ball rolls over a plate of finite dimensions, the interference between the sound generated at the point of impact and the wave reflections at the edges of the plate causes a characteristic spectro-temporal variation of the sound. The goal of the perception experiments reported in this chapter was to investigate whether this acoustic information can be used by listeners to derive information about the rolling direction of a sound. The sounds used in the experiments were either recorded or synthesized and represented balls that were rolling towards or away from the edge of a plate. In the first two second experiments, the participants had to indicate which of two sounds was caused by a ball rolling towards the edge, and which was caused by a ball rolling away from the edge. No training was provided and no feedback was given about the correctness of their responses during the measurement. The overall performance was close to chance. In a control experiment it was established that participants could easily *discriminate* between sounds that differed only in the rolling direction, despite the fact that they did not receive any training with these sounds. Thus, although the sounds contained enough auditory information to be perceived as different, the participants were not able to interpret these differences consistently in terms of the rolling direction.

5.1 Introduction

In the preceding chapter we examined the acoustic effect of the finite dimensions of the plate on the vibrations caused by a rolling ball. More specifically, we argued that a ball rolling towards an edge of the plate can be modeled by extending the

plate beyond the edge, and rolling an imaginary mirror image of the ball towards the same edge, but in the opposite direction of the real ball, over this imaginary extension of the plate. The vibrational waves take some time to travel from the imaginary part of the plate to the real part of the plate, and only the real part of the plate can radiate sound. Therefore, the acoustical effect of the mirror image of the ball is a delayed version of the rolling sound. In the frequency domain this results in comb filtering of the sound. When the ball approaches the edge, the amount of delay becomes less, and thus the distance between the crests of the comb filter becomes larger. This is thus a spectro-temporal effect that varies with the distance of the ball to the edge of the plate. Hence, this effect could be used by listeners to determine the position of the ball on the plate.

The varying comb filter effect is influenced by many parameters of the plate, such as the thickness and the plate size, and the speed of the ball, but not by any of the geometrical or material properties of the ball. It might be less audible when the ball bounces heavily but, in essence, it is not altered. In other words, it conveys information about the ball position, and thus speed, but is invariant to many other parameters, which could make it a good source of information for the perception of the speed of the ball. Because of the simple relation between position and speed, listeners could use this spectro-temporal effect also for estimating the rolling speed of the ball. We will confirm in Chapter 7 of this thesis, Houben's findings [40] that listeners are generally well capable of identifying the faster of two rolling balls.

In other words, we want to test if people can use this spectro-temporal effect to determine the speed of the ball. If listeners can interpret information that is presented to them via this spectro-temporal effect, and use it to determine the speed of the ball, they should be able to determine the movement of the crests of the comb filter. If listeners determine this movement, they should also be able to note if the ball moves towards the center or towards the edge of the plate, because in the first case the crests of the comb filter move downwards in frequency, and in the second case they move upwards in frequency, as shown in Figure 4.10. By asking for this rolling direction, we exclude other cues normally used by listeners to determine the rolling speed of a ball, such as the spectrum and temporal effects due to the imperfect sphericity of the ball, simply because those are independent of rolling direction.

5.2 Experiment

Two experiments are carried out to test whether listeners can hear if a ball rolls from the edge towards the middle of the plate or starting in the middle and approaching an edge of the plate. In experiment 1 we used recorded sounds, in experiment 2 we used synthesized sounds. The synthesized sounds, used in the second experiment, sound cleaner and the comb filter effect is better audible. This is due to the simplified synthesis algorithm and the absence of all other temporal variations.

5.3 Synthesis

For the synthesized sounds, a procedure was used where the average spectrum of the corresponding ball was copied onto a noise signal, whereafter the effect of the reflected waves was added, as discussed in the preceding chapter and by Houben and Stoelinga [44].

From measurements we know that the spectra of rolling balls exhibit a band-pass characteristic of which the shape depends on various properties of the ball, such as its velocity. One possibility to synthesize a rolling sound is by using a physical model which describes the interaction between ball and plate. Such a model is discussed in Chapters 6 of this thesis, but it is not yet capable of synthesizing convincing rolling sounds for the case that the ball-plate contact is continuous. On the other hand, the effect of the plate edge is easier heard in the sound when there is a continuous contact between the ball and the plate. Therefore, we chose to use steel balls, which give a continuous ball-plate contact, and to use a different synthesis algorithm. The model only requires the basic geometric configuration measures of the ball and plate, and the average spectrum of a rolling sound.

The simplest model for generating this non-varying sound consists of band-pass-filtered noise. This results, however, in an unrealistic, synthetic sounding stimulus. A better result, which mimics more precisely the spectrum of a rolling sound, is obtained by taking a sample of white noise, and filtering this white noise with a special filter that gives the noise sample the same, long-term average spectrum as a chosen rolling sound. The temporal properties of the sound are due to the finite plate dimensions and the imperfect sphericity of the balls. By taking the long-term average spectrum of the sound these temporal properties are not represented. The long-term average spectrum is calculated by first splitting the

recorded sound into frames, then determining the spectrum of each frame and, finally, averaging the spectra of these frames. Half overlapping and windowed frames were used to avoid artifacts due to the edges of the frames. The frames must be long enough to cover the lowest frequencies of the rolling sound. This lower limit is determined by the lowest frequencies radiated by the plate. To be on the safe side, hanning windows of 1024 samples were used, which results for a sample frequency of 48 kHz in a lowest possible frequency of 46.9 Hz.

Vibrational waves travel from the point of excitation, the ball, to the side of the plate, where they are reflected by the edge of the plate. These reflected waves appear to come from a virtual source outside the plate. In the case of a string, those mirror images are easily determined mathematically. In the case of a plate, however, this is mathematically not obvious due to added complexity by the fourth-order differential equation and the fact that a plate is modeled two-dimensionally instead of one-dimensionally.

We only added the effect of the first reflection of the vibrations because this one is stronger than the other reflections later on, since the vibrations are damped when they travel across the plate. The sound radiated from the plate due to vibrations from the mirror image occurs somewhat later than the sound radiated due to the direct vibrations of the ball. The effect of a time-varying delay can be approximated by splitting the sound into frames and varying the delay between frames. We again used half overlapping windows with a length of 1024 samples, using a sample frequency of 48 kHz. The delays, ranging from 0 to 4 ms, had very small steps between two consecutive segments, so that the steps in the delay time are not heard. Two complicating effects have to be included to obtain a natural sound. The velocity of the vibrational waves is not constant but depends on the square root of the frequency. A frequency-dependent delay was synthesized by calculating one period of a Schroeder phase complex [89], and this phase complex was used to obtain the coefficients for a FIR filter. A Schroeder phase complex is a sound in which all harmonics are of equal amplitude,

$$r(t) = \sum_{n=1}^N \cos\left(\frac{2\pi n f_0 t}{T} + \phi(n)\right), \quad (5.1)$$

and have a phase relation $\phi(n)$ as follows:

$$\phi(n) = \frac{2\pi C n(n-1)}{N}, \quad (5.2)$$

where n is the harmonic number and N the total number of harmonics. C was chosen to be 1, in which case the tone complex is called Schroeder positive. The signal approaches a periodic, linear FM chirp. The instantaneous frequency decreases linearly from the frequency of the highest component, which was 24 kHz, down to the fundamental frequency, 250 Hz for the longest delay.

This range of frequencies is covered in maximally 4 ms, making it a rather fast chirp. A second effect of the propagation of vibrational waves in a plate is the faster damping of the higher frequency vibrations. This was modeled by multiplying the time signal with a ramp. The chirp starts with a high frequency and therefore only those high frequencies are attenuated by the ramp. This damping as well as the delay times are a property of the material and thickness of the plate. These values should therefore be altered when synthesizing the sound produced by other plates.

5.4 Method

The experiment took place in a quiet room, where the participants were presented with the stimuli via headphones. The experimental paradigm was two-alternative forced-choice. Feedback was never given. After given an introductory text, participants were three times presented with a set of 18 stimuli. Each set consisted of the same 18 pairs of sounds, presented in a different random order. The participants were asked to respond to the question “Which ball rolls towards the edge of the plate?”, and presented two buttons labeled “first” and “second” to indicate their answer. Feedback about the correctness of their answers was never given.

This relatively short experiment took the participants about five minutes. They were not paid for participating. The participants were colleagues who were aware of the general line of research but uninformed about the effect of the plate edge on the sound of rolling balls. In both experiments we had four participants. None of the participants that volunteered in the first experiments also participated in the second experiment.

5.5 Stimuli

The sounds were recorded with the setup described in the preceding chapter. Steel balls with diameters of 20, 35 and 50 mm were used, which were rolled at speeds of 0.67, 1.15 or 1.6 m/s. These sounds were also used to obtain the

spectrum for the synthesized sounds. From all combinations two samples were taken, one sample was taken from the beginning of the sound where the ball rolls from the edge towards the middle of the plate and another sample starting about halfway the sound where the ball rolls from the center towards the opposite edge of the plate. Both orders of these two samples formed one stimulus. Due to the different speeds of the balls, the samples had different lengths, ranging from 500 to 900 ms. In between the stimuli there was always 500 ms of silence. The total number of stimuli was $3 \times 3 \times 2 = 18$.

TABLE 5.1: Results of experiment 1, using recorded sounds. In the middle part the percentages correct of the responses are given for the four participants and the three sets of stimuli, which represent the same 18 stimuli in different random orders. On the lowest row the mean percentages correct for each participant are presented, on the right side the mean percentages for each set are presented, while the grand mean of all responses is given at the bottom right.

		participant number				
		1	2	3	4	mean
stimulus set number	1	44	39	72	72	57
	2	72	56	50	44	56
	3	39	44	56	61	50
mean		52	46	59	59	54

TABLE 5.2: Results of experiment 2, using synthesized sounds. In the middle part the percentages correct of the responses are given for the four participants and the three sets of stimuli, which represent the same 18 stimuli in different random orders. On the lowest row the mean percentages correct for each participant are presented, on the right side the mean percentages for each set are presented, while the grand mean of all responses is given at the bottom right.

		participant number				
		5	6	7	8	mean
stimulus set number	1	39	44	44	50	44
	2	56	44	44	56	50
	3	56	44	33	44	44
mean		50	44	41	50	46

5.6 Results

The experimental results show that listeners cannot or, not very well, detect the rolling direction of a rolling ball on the plate, which was either rolling from or towards the edge of the plate. For each of the two experiments, the two-sided, 95% confidence interval for guessing covered the range from 43.5% to 56.5%. The mean percentage correct values of both experiments fall within this region. When looking at each participant individually, the two-sided, 95% confidence interval for guessing covered the range from 37.0% to 63.0%. Here too, all results fall within this region. There is no trend in the percentages correct for consecutive sets, and thus we have no indication that the participants are learning their task. We can thus not conclude that the participants extracted any relevant information about the rolling direction from listening to the sound.

5.7 Control experiment

The negative result in the first two experiments described in this chapter could have different reasons. One explanation is that listeners can *discriminate* between sounds that roll towards or away from the edge, but that they are unable to *identify* which sound represents rolling towards the edge. Another more trivial explanation is that listeners cannot *discriminate* between two sounds that differ in rolling direction.

To test this latter explanation, a control experiment was set up. The same stimuli were used as in the first two experiments described in this chapter. To begin with, two sounds were presented to the listener as before. They again originated from three different balls rolling with three different speeds. These sounds were presented in two versions, one approaching the edge of the plate and the other moving away from the edge of the plate. After these two sounds a third sound was presented, being equal to one of the first two sounds. The participants' task was to indicate whether they thought this third sound was equal to the first or second sound. The number of stimuli was 36 sound triplets per set this number is twice as large compared to the first two experiments, because there are two alternatives for the third sound in the triplet. This time all the experiments consisted of three sets of recorded sounds and three sets of synthesized sounds for all eight participants. The sets were presented in alternating sequence, beginning with the recorded sounds. Three of the participants were also participants of one

of the first two experiments. This control experiment took the participants about 15 minutes, they were not financially rewarded.

5.8 Results

TABLE 5.3: Results of the control experiment. In the middle part the percentages correct of the responses are given for the eight participants and the two types of stimuli, either recorded or synthesized sounds, averaged across all three stimulus sets. On the lowest row the mean percentages correct for each participant are presented, on the right side the mean percentages for each stimulus type are presented, while the grand mean of all responses is given at the bottom right. Numbers presented in italics are not different from guessing in a statistically significant way.

	participant number								mean
	1	2	3	4	5	6	7	8	
recorded sounds	91	96	86	100	99	83	91	83	91
synthetic sounds	76	61	50	56	96	70	77	81	71
mean	83	79	68	78	98	77	84	82	81

For the recorded sounds, the participants could clearly distinguish the two sounds, with the worst score still 83% correct. For synthetic sounds, the situation is somewhat different. The individual differences were quite large. Two out of eight participants were not able to score significantly better than chance. Another participant scored better than chance in a statistically significant way ($p=0.011$), but he or she had difficulties in doing the task, scoring only 61% correct. Only five out of eight participants scored reasonably well, of whom one scores very well. The participant that scored equal to chance (50% correct), indicated afterwards that he or she could not do the task for these sounds and continued by always pressing the same button.

5.9 Discussion

In the previous chapter it was explained that the wave reflections result in peaks in the spectrum. These peaks move upwards in frequency when the ball ap-

proaches the edge of the plate and downward in frequency when the ball moves away from the edge of the plate. In the third experiment of this chapter, it was shown that, for the recorded sounds, and for five out of eight participants also for synthetic sounds, listeners could very well *distinguish* between the sound from a ball rolling towards the edge of the plate and one rolling towards the edge of a plate. In an *identification* task, both experiments showed performance that was not significantly different from chance. This means that the listeners could not tell where the differences they heard originated from.

The time needed for the ball to reach an edge of the plate can also be derived from the movement of the peaks in the spectrum, in a similar way that plummeting gannets estimate the time before they hit the water [59], and this could help listeners in estimating the speed of the ball. With the current negative results regarding the rolling direction, it seems unlikely that listeners do use this spectro-temporal pattern for this seemingly more difficult task. Thus, the acoustic information resulting from the reflections at the plate edge does not help in explaining the generally good capability of listeners to discriminate the faster rolling ball from the slower one, which will be shown in Chapter 7.

6 | Time-domain Modeling and Simulation of Rolling Objects¹

Abstract

A model that was previously developed for simulating a single impact of a ball on a damped plate has been adapted to simulate the sound of a ball rolling over a plate. The original model has the advantage of being well tested and showing good agreement between measurements and time-domain simulations of various impacted plates [12]. The main changes for its adaptation to rolling sounds were made in the ball-plate contact. Instead of an impact point that is fixed in space and short in time the model now incorporates an interacting contact point that is continuously moving in space. To allow for the variable position of the contact point, we use a special spatial window that is optimized for this purpose. Furthermore, a model for the surface roughness of the plate was added.

The model is validated by means of three different types of simulations. The numerical results are either compared with experiments or with analytical calculations. The first type of simulation is that of a ball that rolls over a surface with some random asperities. The main observation is that the ball loses contact with the plate at some speed. The second simulation is that of a sinusoidally time-varying source that moves over the plate. Here the characteristic Doppler effect is identified. The third set of simulations are of a ball that is dropped on a plate. The ball bounces back to some height that is lower than the original release height. This fraction of height, also called the *restitution coefficient*, was measured and compared with simulated data.

Following the validation procedure, the model is used to simulate rolling objects. It is shown that different kinds of contact exist between ball and plate. Four

¹This chapter is based on Stoelinga and Chaigne [97] and [93]

different types of rolling with different plate/ball contact parameters are identified: amplitude modulations, periodic bouncing, chaotic bouncing and continuous contact. Comparisons are made between measured and simulated accelerations of a fixed point on a aluminum plate with a sinusoidal waviness profile, which is set into vibrations by rolling spheres of various sizes, stiffnesses and densities.

6.1 Introduction

Many attempts have been made to synthesize rolling sounds. These roughly fall in two categories. They can be based on signal processing or based on numerical simulation of physical laws. The signal processing models can also have some physical basis. Often we can distinguish several steps in the processing, representing a surface model, a contact model and a plate model.

The signal based approach was used by Hermes [36], who tried to generate the sounds by convolving a series of impulses with a gamma tone, a cosine manipulated to have an envelope formed by a polynomial attack and exponential decay. Some temporal variations are added (amplitude modulation and a decaying amplitude of the pulses). Also Van den Doel *et al.* [105] used this approach, their focus lied on the generation of sounds for interactive virtual environments. One example they give is the rolling and bouncing of a rock in a virtual wok. Pauly *et al.* had the same focus, and presented a model including a two-dimensional contact including a Coulomb friction model [76]. They were capable of simulating complex situations such as a spinning ball dropped onto an inclined plane.

Within the sounding object project, an attempt has been made to synthesize rolling sounds [81]. The researchers from this project made, as they called it, “cartoonifications” of a real rolling sound. The strong point of their algorithm is the fact that the sounds are generated in real time, allowing their use in virtual reality applications. On the other hand it is impossible to compare experimental recordings with computed sounds, because the latter are only intended as a cartoon of the first. The only information about the quality of their sounds is available from the effectiveness of the sounds in an interactive augmented reality experiment [80].

For train track-wheel interaction a lot of research has been done, and the models have a physical basis. One of these numerical models, called TWINS, was verified with experimental data [101] [100]. The authors conclude that “The model is found to give reliable results, provided that extreme care is taken with the input data, especially concerning the roughness data and the attenuation of the vibration along the rail”. The model, however, is defined in the frequency domain and, due to the lack of phase information, we cannot transform it to the time domain and listen to the results. Furthermore their model is linear, and thus it is expected to fail in simulating the chaotic bouncing phenomena observed in experiments, as described in Section 6.4.3 of this paper. A non-linear time domain model was proposed by Nordborg [74]. The output of this model, however, was not compared to measured data in his paper.

The non-linear interaction between a sphere and plane has been subject of research also. Some interesting results have been found in this way. In one of the first articles on the topic, by Nayak [72], the loss of contact due to vibrations of the plate is calculated, using the method of harmonic balance. Later, similar results were found with the more refined method of multiple scales [38].

By applying the stationary solution of the Fokker-Planck equation to the Hertzian contact problem, the resulting stationary probability density function for the relative displacement can be derived [75]. The interaction of two rough surfaces as well as surface roughness of differently finished surfaces were presented by Sayles and Soom [87].

As a basis for the current work we start from the model of impacted plates by Chaigne and Lambourg [12] [57] which is summarized briefly in Section 6.2. It has the advantage of yielding a good agreement between measurements and simulations of damped impacted plates. This agreement is primarily due to the detailed modeling of the various processes contributing to the damping in the plate. This model consists of three parts. One part is the contact model between the ball and the plate based on Hertz’s law. Since it had to be changed here to allow for rolling sounds, this adaptation is treated in Section 6.3. The contact of the ball sets the plate into vibration. The second part of the model is that of the flexural vibrations in the plate. In the numerical algorithm, the plate is represented at discrete gridpoints. In all these gridpoints the motion of the plate is calculated by solving the Kirchhoff-Love plate equation using an explicit finite difference scheme.

The schemes used are of 4th order in time and 2nd order in space. An important aspect for the sonic quality of the simulations is the damping in the plate. In the original model thermoelastic damping, viscoelastic damping, and damping due to the radiation of sound are accounted for. The third part of the model is a model for the radiation of the impacted plate sound in free space. This is covered by calculating a Rayleigh integral over the plate. A simple correction at low frequencies is applied to account for the case of an unbaffled plate. This radiation model has shown to be effective and was not changed. In Section 6.4, preliminary simulations of spheres rolling on an MDF (Medium Density Fiberboard) plate are presented in order to validate the model. The ability of the method to reproduce typical features such as Doppler effect, restitution coefficient at the rebounds, and the influence of surface profile on the particular histories of rolling sounds are presented and discussed. Finally, Section 6.5 presents an extensive comparison between measured and simulated plate accelerations for an aluminum plate with a specially designed sinusoidal profile, which is set into vibrations by rolling spheres of various sizes, stiffnesses and densities.

6.2 Single impact on a damped plate

We consider here the sound radiated by the flexural vibrations of a finite plate excited by a single impact. This excitation will be revisited in the next section to account for rolling. To simplify the presentation, it is assumed that the plate is rectangular, thin, homogeneous and isotropic, and that the Kirchhoff-Love approximations are applicable. The thickness h of the plate is a function of the coordinates (x, y) , to allow for introducing roughness and waviness in the model. In the simulations, the boundary conditions can be selected either free, clamped or simply supported on each edge. In the experiments presented in Sections 6.4 and 6.5, the plates are free on each edge and are lying on four springs: this has the consequence of introducing a supplementary rigid body mode of very low frequency with no consequence on the present problem, so that it will not be discussed any longer in this paper.

One main challenge of the present study is to reproduce rolling sounds and vibrations in the time-domain convincingly so that the results of the simulations can be heard, and reasonably compared to real recorded sounds. This imposes to use an accurate model of damping for the plate. We use here a general model of damping that includes thermoelastic, viscoelastic and radiation losses [57]. In what follows, two different isotropic plates were used: MDF and aluminum, re-




FIGURE 6.1: Geometry of the plate.

spectively. The MDF plate is characterized by high attenuation, which is well represented by a viscoelastic model. For the aluminum plate (as for metallic plates, in general), only thermoelastic and radiation losses are taken into account. As a consequence, the rigidity of the plates becomes complex. In what follows, the derivatives versus time $\partial/\partial t$ are replaced, for convenience, by a multiplication by s , the Laplace variable. For the MDF plate, the complex rigidity is given by [12]:

$$\tilde{D} = \frac{Eh^3}{12(1-\nu^2)} \left[1 + \frac{sR_{v1}}{s + s_{v1}} + \frac{sR_{v2}}{s + s_{v2}} \right], \quad (6.1)$$

where E is the Young's modulus of the plate and ν its Poisson's ratio. R_{v1} , R_{v2} , s_{v1} and s_{v2} are viscoelastic parameters derived from curve-fitting of measured decay times of free vibrations over a broad range of frequencies [57]. The magnitudes of both fractions between the brackets in Equation (6.1) are supposed to be small expressed relative to unity. The numerical values used here are listed in Table 6.1 on page 97. For the aluminum plate, the complex rigidity is written [12]:

$$\tilde{D} = \frac{Eh^3}{12(1-\nu^2)} \left[1 + \frac{sR_t}{s + c_t/h^2} + \frac{2\rho_a c_a}{\omega_c \rho h} \frac{\sum_{m=1}^3 b_m (s/\omega_c)^m}{\sum_{n=0}^3 a_n (s/\omega_c)^n} \right], \quad (6.2)$$

where ρ is the plate density, R_t and c_t are thermoelastic constants of the material, ω_c its critical frequency, ρ_a is the density of air, c_a the speed of sound in air, and a_n and b_m are constants (see Table 6.2). The two fractions between the brackets in Equation (6.2) are perturbations terms of the first order compared to unity. The first fraction accounts for thermoelastic damping while the second fraction accounts for radiation losses in the plate and takes the form of a Padé approximant.




FIGURE 6.2: Ball-plate interaction. Hertz's model.

Numerical values of these physical parameters can be found in Table 6.2.

As a consequence of this damping model, the transverse displacement W of the plate is governed by the equation:

$$\rho h s^2 \tilde{W} = -\tilde{D} \nabla^4 \tilde{W} + \tilde{q}, \quad (6.3)$$

where $\tilde{W}$ and $\tilde{q}$ are the Laplace transforms of displacement and excitation, respectively. ∇^4 is the biharmonic operator.

It is assumed that both plate and impactor are elastic solids and that their interaction is governed by Hertz's law of contact. In this context the interaction force at the contact point of coordinates (x_0, y_0) between the plate and the sphere of radius R is given by (see Figure 6.2):

$$F_c(t) = \begin{cases} \kappa [R - (\eta(t) - W(x_0, y_0, t))]^{3/2} & \text{for } \eta(t) - W(x_0, y_0, t) < R \\ 0 & \text{otherwise,} \end{cases} \quad (6.4)$$

where $\eta(t)$ is the displacement of the center of gravity of the impacting sphere (or ball) of mass m . κ is the Hertzian constant, in $\text{N}/\text{m}^{2/3}$ [107]. During the contact, the motion of the sphere is governed by Newton's second law:

$$m \frac{d^2 \eta}{dt^2} = -F_c(t). \quad (6.5)$$

The origin of time $t = 0$ is selected at the very instant when the sphere just hits the plate with initial velocity $\frac{d\eta}{dt} = V_0$.

In Equation (6.3), q has the dimension of a pressure. In order to establish the link between the impact force and q , one has to consider that, during the contact phase, the force is distributed over a finite area. To account for this, we define a spatial window $M(x - x_0, y - y_0)$ so that:

$$q(x, y, t) = M(x - x_0, y - y_0)F_c(t). \quad (6.6)$$

The order of magnitude for M was obtained experimentally by measuring the spot on plates consecutive to the impact of spheres previously colored with ink. In addition, the model allows temporal variation of M to account for the variation of the contact surface with time. However, no significant discrepancies were observed between simulations and measurements, at least for hard impactors [11].

Equations (6.3)-(6.6) are solved numerically by means of high-order explicit finite difference schemes [57]. In order to minimize frequency warping, a high sampling frequency f_e is requested: 192 kHz was the standard value used, which ensures estimation of the eigenfrequencies of the plate up to 5 kHz within one per cent accuracy. However, for appropriate reproduction of the restitution coefficient (see Section 6.4.2) it was found that higher sampling rates were necessary.

The numerical model yields the time history of impact force $F_c(t)$, plate displacement $W(x, y, t)$ and sphere displacement $\eta(t)$ as output files. From the surface distribution $W(x, y, t)$, the radiated sound pressure $p(\mathbf{r}, t)$ at a given point $\mathbf{r}$ in space is computed, using a Rayleigh integral [57]. Other vibrational quantities, such as plate velocity and acceleration, can also easily be derived.

6.3 Adaptation of the model for moving sources

We start the explanation of the model for rolling objects by considering first “imposed” moving sources, thus leaving the plate-ball interaction problem temporarily aside. This preliminary simple model is used to explain the distribution of force onto the gridpoints in section 6.3.1, and it will also be used later in section 6.4.1 to show vibrations in the plate in the case of a moving, harmonic excitation. An imposed force source would correspond roughly to the case of a shaker

attached to the plate. However, for a moving source such an attachment is not possible in reality, and thus this model is rather theoretical.

In a second part, a more realistic case is treated where the vibrations of the plate lead to a change of force between ball and plate. This model is called “interactive”. This model was used in almost all the simulations, except in the previously mentioned ones.

6.3.1 Windowing

In our numerical model, the plate is represented at discrete positions, in the time domain as well as spatially. The discrete spatial locations where the displacement of the plate is defined will be called *gridpoints* hereafter.

To move a force smoothly over the plate, we need to convert the contact force, F_c , located at a continuously defined position to one at the discrete gridpoints, (N_x, N_y) , of the plate. Thus, a numerical version of Equation (6.6) is now written:

$$q(N_x, N_y) = F_c \times M(N_x \Delta x - x_0, N_y \Delta y - y_0). \quad (6.7)$$

The contact force is spread a little in space, thus over the gridpoints, by the window function $M(\dots)$. The gridpoints are separated by a spatial step $\Delta x, \Delta y$. In order to create a window in two dimensions, as needed by Equation (6.7), we multiply a window in the two orthogonal directions:

$$M(x, y) = M(x)M(y). \quad (6.8)$$

The windowing is illustrated in Figure 6.3. For moving sources, the appropriate choice of the window function is more critical than for fixed impacts. To satisfy numerical accuracy, some constraints are posed on the window function. One first constraint is that, for a given imposed force, this force is kept constant when the point of impact is changing. The second constraint is that the windowing should be independent of the spatial sampling. These constraints are always met by the class of generalized cosine windows, and therefore these functions were selected (see Figure 6.4). The generalized cosine windows have the form




FIGURE 6.3: The spatial window spreads the force onto the gridpoints. When the ball rolls over the virtual plate, the position of the window changes. Due to this effect at each point in time, a different portion of the window is found at the gridpoints. The spatial window is thus converted into a window in time, that is stretched for higher velocities.




FIGURE 6.4: Plot of the various window functions. The upper left panel shows the window function itself, the upper right panel shows the second derivative and the lower panel shows the fourth derivative of the window function. The width of the windows is normalised to their minimal usable width. Although the functions itself look much alike, their higher order derivatives are very different. The s4-window (see Equation 6.12) is the only one that is zero at its border at all derivatives

$$M(x) = \begin{cases} \sum_{n=0}^N a_n \cos(n\pi x/2) & \text{if } -1 \leq x \leq 1, \\ 0 & \text{(otherwise),} \end{cases} \quad (6.9)$$

where N is the order of the window function. Using $a_0 = a_1 = \frac{1}{2}$, we obtain the Hanning window. Generally the window is normalized so that $M(0) = 1$. In our case, we wish that the window be smooth near the edge, in its first, second, third and fourth derivatives, in order to avoid unwanted noise in the simulations due to unrealistic discontinuities. This implies:

$$M(0) = 1, M(1) = 0, M''(1) = 0, M'''(1) = 0, \quad (6.10)$$

for which we need four degrees of freedom, or $N = 3$, and this leads to

$$a_0 = \frac{5}{16}, a_1 = \frac{15}{32}, a_2 = \frac{3}{16}, a_3 = \frac{1}{32}, \quad (6.11)$$

or in combining with Equation (6.9):

$$M_{s4}(x) = 5/16 + 15/32 \cos(\pi x/2) + 3/16 \cos(2\pi x/2) + 1/32 \cos(3\pi x/2). \quad (6.12)$$

This last function will be called the smooth-fourth-order derivative generalised cosine window or abbreviated “s4-window”. It is also possible to generate a window function that is only smooth in its first two derivatives. This function,

$$M_{s2}(x) = 3/8 + 1/2 \cos(\pi x/2) + 1/8 \cos(2\pi x/2), \quad (6.13)$$

we will call “s2-window”. The relevance of the appropriate choice of windows is shown in Figure 6.5. We can see that the numerically calculated second-order derivative of the hanning window has large errors near its discontinuities which do not vanish when the spatial discretisation step is taken smaller. The other window shows much less error, and this error does decrease with decreasing spatial step. The error depends largely on the position of the window, that is, on the position of the highest point of the window in between the gridpoints. This is shown in Figure 6.6. If the window is approximately centered, the error is large. The error is periodic with the sample interval: two periods are shown in Figure 6.6. Note that in Figure 6.6 the error values are not reduced when using more sample points in the case of the hanning window, because the largest source of error is the discontinuity. For the s2-window, on the contrary, we can reduce the error as




FIGURE 6.5: A window function with its second-order derivative calculated twice, once analytically and once numerically. We see in the left pane a Hanning window with a large difference between the two double derivatives, due to its discontinuous second order derivative. In the right pane, a generalised cosine window with smooth second order derivative is depicted, which has no large discrepancy between the analytically and numerically calculated second-order derivative.

much as we wish by taking more sample points.

For the sake of simplicity and brevity, only the properties of the “s2-window” are presented in Figures 6.5 and 6.6. In fact, due to the presence of 4th-order derivatives in the plate equation, an “s4-window” was used preferentially in our simulations.

6.3.2 Interactive model

To create an interactive model, where the force depends on the displacement of the plate, we need to know this displacement in the first place. However the plate displacement is only calculated at the gridpoints. To find the displacement at the contact point, it needs to be interpolated from the nearby gridpoints. This is the inverse problem as treated in the previous section, and the same windowing technique is used to obtain the plate displacement at the position of the contact.

We return now to the model used for the contact. As for isolated impacts, see Equation (6.4), the ball-plate interaction is assumed to be governed by Hertz’s law. To account for surface irregularities (random rugosity and/or waviness), we make now the distinction between the plate displacement W_p and the surface profile W_s . This distinction is necessary because the scale of the irregularities is usually smaller than the scale of the spatial grid used for the plate and thus, we




FIGURE 6.6: The actual error of the numerical method largely depends on the shortest distance between the highest point of the window and the discretisation points. Here we see the error in the case of a window sampled at eight points (left) and at sixteen points (right).

will have to process them separately. As a consequence, following the procedure used by several authors [74] [107], the interaction force is rewritten as follows:

$$F_c = \begin{cases} \kappa [R - (\eta(t) - W_p(x, y, t) - W_s(x, y, t))]^{3/2} = \kappa \alpha^{3/2} & \text{for } \alpha > 0 \\ 0 & \text{otherwise,} \end{cases} \quad (6.14)$$

where the coordinates (x, y) of the contact point are now moving with time. Notice that, in some cases, a damping term of the form $\mu \dot{\alpha} \alpha^{1/2}$ was added in the expression of F_c in order to account for dissipation in the ball [47][55].

6.3.3 Surface model

Although for impacts it is common to regard the ball and plate as perfect sphere and plane, in rolling the asperities of the contacts are the source of the structural vibrations. There have been a number of publications dealing with the contact of two rough surfaces. Two different approaches can be identified. One is a stochastic approach, sometimes called a Greenwood and Williams surface model, after the authors of the pioneering work [32]. In this model the two rough surfaces are replaced by a smooth surface in contact with an equivalent rough surface. The equivalent rough surface consists of asperities with simple geometrical shapes. They are assumed to have a stochastic distribution for some parameters such as the asperity height and density [51][52]. In a fractal approach, the two surfaces are considered to have the same self-affine scaling properties, but are uncorrelated [45]. These publications do not treat the case of a curved surface, but if the




FIGURE 6.7: Interactive model of rolling on a rough surface with two spatial scales. The interaction of the sphere with the surface profile is resolved at a *fine* spatial scale. The resulting force is transferred to the plate using a spatial window. The plate displacement is then calculated at its gridpoints (*large* scale). This displacement is then interpolated on a finer grid so that it can be included in the calculation of the contact force. The procedure is repeated at the next time step, for a different position of the moving sphere.

surface irregularities are small compared to the sphere size this approach seems appropriate.

In our numerical simulations, a profile map was used to describe the relative height of the surface. Some stochastic as well as a fractal methods were tried to generate a profile map. These methods gave different simulation results, but, until we have done an extensive study of these differences, we choose to use a very simple model. For the simulations, the relative height of the surface is stored in a profile map, representing the relative height of the plate surface. Before the actual simulation starts, a profile map is filled with random numbers between zero and a given maximal height. This leads to a surface height that is uniformly distributed between zero and this maximum height. This profile map is then used in Equation (6.14) to calculate the contact force.

Usually, the surface irregularities, and the contact region, are defined on a much smaller scale, about 50 times smaller than the size of the plate grid. It is unfeasible to change the grid to such small sizes since it would lead to 2500 times longer calculation times. Therefore we calculate the contact force with a spatial grid that is much finer than the grid used for the calculations of the plate vibrations, and then distribute the force that was found over the nearby points of the plate grid, using the earlier introduced windowing functions. Conversely, the resulting displacement of the plate is interpolated so that the contact force can be

calculated on the fine grid corresponding to the rugosity scale of the surface. This procedure is explained in Figure 6.7.

6.4 Simulations

In this section, a first series of simulations is presented in order to validate the model. In the next paragraph, a simple constant moving force is applied to the plate in order to test the ability of the model to reproduce Doppler effects adequately. In Section 6.4.2, the excitation is fixed and the ball-plate interaction model is tested by comparing measured and simulated restitution coefficients. In Section 6.4.3, the complete interactive model is used and comparisons are made between measurements and simulations of spheres rolling on an aluminum plate with a sinusoidal surface profile. Here, we focus on the critical velocities for which there is a loss of contact. Simulated velocities are compared with measurements performed using a high-speed camera. Finally, large ranges of input physical parameters are tested in Section 6.4.4 in order to explore the capability of the model to simulate a significant variety of sounds: here the sounds and vibrations produced by different spheres rolling on an MDF plate with imposed random rugosity are simulated.

6.4.1 Doppler effect

One effect that is immediately expected when moving sources are involved is the Doppler effect. Due to the speed of the sphere, the vibrations in front of it are compressed while the waves behind are elongated. Figure 6.8 shows an example of simulations for a sphere rolling on a very low and narrow metallic plate (a rail). Observations are made before the waves reach the end of the rail. On this figure, we can see the compressed and expanded waves, as expected. The frequency shift is expected to be asymmetric around the forcing frequency f_f where the higher frequency, in front of the source, equals $f_f/(1+M)$ and the lower frequency, behind the source equals $f_f/(1-M)$. The mach number, M equals 0.207 in our case resulting in 631 and 414 Hz, before and behind the source, where we found 638 and 419 Hz in our simulation.

6.4.2 Restitution coefficient

When a ball is dropped on a plate, it bounces back to some height that is lower than the original release height. The fraction of the original and rebound height




FIGURE 6.8: The plate vibrations at a fixed point vary in frequency due to the Doppler effect. The example shown is obtained for a 500 Hz sinusoidal imposed force moving with a speed of 500 m/s. The origin of time corresponds at the very instant where the source passes at the fixed point. At the left side where the source approaches the observation point, the average frequency is 638 Hz, and it is 419 Hz at the right side.

is called the *restitution coefficient*. The restitution coefficient represents the loss of energy during bouncing. It varies with many material parameters due to various effects, as described in the literature [20][111].

Simulations are made here on MDF plates, which are also used for the experiments. Parameters of these plates are found in Table 6.1. The predominant energy loss is due to the spread of vibrational energy in the plate. As seen in Figure 6.9, our model gives good results for the plate with thickness $h = 6$ mm, whereas the energy loss is underestimated in the model by nearly 20 % for the thick plate ($h = 12$ mm). This might be due to the fact that the model is based on the thin plate theory, which is probably not justified anymore in this latter case.

Finally, it might be of interest to indicate that a large sampling rate was required in order to estimate the restitution coefficient with sufficient accuracy. Figure 6.10 shows an example of this, for a ball bouncing on the thinnest plate ($h = 6$ mm), where it is shown that the estimation of this coefficient converges to an asymptotic value for sampling frequencies above 1536 kHz.




FIGURE 6.9: Restitution coefficient of a ball bouncing on a plate for different plate thicknesses and impact velocities. The circles represent measured data for a 10 mm steel ball bouncing on a 12 mm thick MDF plate and the crosses are for the same ball bouncing on a 6 mm thick MDF plate. The lines represent simulated data for the two cases.




FIGURE 6.10: When the sampling frequency is too low, the simulated restitution coefficient is affected by numerical errors. Estimation of this coefficient stabilizes after 1536 kHz.

6.4.3 Loss of contact and rebounds

When a ball rolls over a surface with sufficient asperities, it loses contact with the plate at some speed. To verify the simulated accuracy of the phenomenon, we rolled a ball over an aluminum plate with a sinusoidal surface profile. The size of the plate was 1×0.2 m with an average thickness of 8 mm. The wavelength of the profile was $\lambda = 5$ cm, and the wave amplitude was $H_0 = 2$ mm. When the ball moves faster over the plate, this profile slides faster underneath the ball. If contact is maintained, the ball is forced to follow the profile, going up and down in vertical direction. At a certain point, the downward acceleration that the ball is required to have to follow the profile exceeds the gravitational acceleration of the ball. This condition for which the ball leaves the plate is given by the so-called “critical velocity” v_0 . A rough estimate of v_0 can be simply obtained by computing this, horizontal, velocity for which the vertical ball acceleration is just equal to the gravity acceleration g which yields:

$$v_0 = \frac{\lambda}{2\pi} \sqrt{\frac{g}{H_0}}. \quad (6.15)$$

For the aluminum plate with wavelength λ equal to 50 mm and amplitude H_0 equal to 2 mm, Equation (6.15) yields $v_0 = 0.56$ m/s. Of course, the main weak point of this estimation is that it does not take the motion of the plate into account. This can be done only numerically and it is in fact one possible application of our model. For the same parameters, and taking further the parameters of the aluminum plate into account, the model yields $v_0 = 0.4$ m/s.

In order to check this value experimentally, we made movies of a ball rolling over the profiled plate using a high speed camera, and the vibrations of the plate were measured with an accelerometer. The exact critical velocity, at which the transition between continuous contact and periodic bouncing occurs, is difficult to determine experimentally. The reason for this difficulty arises from the accelerometer signal, since the ball impacts are difficult to be distinguished from the periodical modulations of the contact force, which are present just before the loss of contact. Even in the images from the high speed camera it is difficult to see a very narrow space between the ball and the plate (see Figure 6.11). In order to circumvent the problem, systematic measurements and simulations were made where the initial speed of the ball was increased stepwise from 0.1 up to 1 m/s with 0.1 m/s steps in order to obtain upper and lower bounds for v_0 .




FIGURE 6.11: Pictures of the highspeed camera. A small ball is rolling over a plate with a sinusoidal profile. The plate is 100 cm long but only the first 30 cm are shown. In the left panel the ball has a velocity of about 0.7 m/s and is above the plate. In the right panel the velocity is about 0.3 m/s and the ball remains in contact with the plate. Note that we see a reflection of the ball in the plate, which can be used to estimate the distance between ball and plate. The complete videos are available at the web address: www.ensta.fr/~chaigne/Rolling_sounds.

In this section we investigated the system without a Hertzian spring. This Hertzian spring can, together with the mass of the ball form a resonant system. Oscillations in this system can also cause loss of contact, and this could be due to very small roughness of the plate the ball rolls over. This effect is analyzed in Appendix 6.C.

For the lowest velocities, below $v_0=0.4$ m/s, there is clearly continuous contact between ball and plate. Around 0.6 m/s we could see the ball impacting the plate. For velocities in between those two values, it is very difficult to determine the actual contact because, at the point where the ball and plate are in slight contact, very few vibrations are induced into the plate. When there is slight loss of contact, the gap between ball and plate is very small making it hard to detect in the films. The results of simulations are summarized in Figure 6.12 where the discontinuity between the two slopes at $v_0 = 0.4$ m/s is in agreement with the experimental observations.

6.4.4 First simulations of rolling sounds

We now use our model to simulate broad categories of rolling objects. We will show that there are different kinds of contacts between ball and plate, depending on their respective geometry and material properties. This list should not be considered as a complete list of all possible rolling types, but rather as typical illustrations. The following examples simulate sounds produced by spheres rolling over an MDF plate whose input parameters are listed in Table 1.




FIGURE 6.12: Average height of the ball above the surface for a ball rolling over a sinusoidally curved surface. Up to a certain critical speed, here 0.4 m/s, the ball remains in contact with the plate. Above this speed the bouncing becomes quickly chaotic.

TABLE 6.1: Input parameters for the simulations with MDF plates

Plate geometry	$1.2 \times 0.5 \times 0.018$ m
Plate material	MDF
Boundary conditions	free
Plate density	700 kg m^{-3}
Young's modulus	$E = 3.4 \text{ GPa}$
Poisson's ratio	0.4
Damping parameters (see [57])	$R_{v1} = 60 \times 10^{-3}$; $R_{v2} = 20 \times 10^{-6}$ $s_{v1} = 15 \times 10^3 \text{ rad s}^{-1}$; $s_{v2} = 50 \times 10^3 \text{ rad s}^{-1}$
Ball velocity	0.6 m s^{-1}
Surface profile	uniformly distributed between 0.0 and 0.1 mm spatial step: 0.19×0.19 mm
Specific parameters in simulations a) to d); See Figure 6.13	
Hertzian constant	a) 80 ; b) 8×10^6 ; c) 8×10^5 ; d) $8 \times 10^4 \text{ Nm}^{-1.5}$
Ball mass (in g)	a) 3.3 ; b) 3.3 ; c) 10 ; d) 20




FIGURE 6.13: Different effects in rolling sounds due to different contact parameters and mass when simulated by the model (see Table 6.1). Left are the simulated acoustical pressures (sounds) and to the right are the contact forces. From top to bottom the four effects are: (a) Amplitude modulations, (b) periodic bouncing, (c) chaotic bouncing and (d) continuous contact. The corresponding sounds are available at the web address: www.ensta.fr/~chaigne/Rolling_sounds.

The first of our rolling types is dominated by an *amplitude modulation* of the contact force. This modulation is due to the resonance frequency of the equivalent contact stiffness and the mass m of the ball. There is no loss of contact here, and therefore the effect is that the vibrations due to the surface roughness are modulated in amplitude. The modulation period T can be calculated, as shown by Falcon *et al.* [20]:

$$T \approx 5.38 \left(\frac{m}{\kappa g^{1/2}} \right)^{1/3}, \quad (6.16)$$

which, in our case, results in $T = 0.127$ s for $m = 3.3$ g and $\kappa = 80$ Nm^{-1.5}, which is in perfect agreement with the period of the oscillations in the simulations (see example (a)) in Figure 6.13.

If the amplitude of the ball motion becomes larger, the ball loses contact with the plate in a periodic manner (example (b)). As it is seen in the second right panel of Figure 6.13, the contacts will be rather short. The bounce height is fixed for a given configuration of ball and plate. It seems to bounce always relatively high, that is, we can hear the bounces on the plate.

For situations where the surface roughness induces the contact loss between ball and the plate, the contact loss is more chaotic (see example (c)). In contrast to the prior situation, the bouncing is not periodic. Also, the bounce height can vary largely. For situations where the contact loss is short, the plate vibrations do not decrease in amplitude during this loss of contact. This we call micro bouncing because in the sound we can not always recognize this bouncing.

The last simulated example, (d), is for a continuous contact between ball and plate. Here the contact acts as a non-linear filter of the surface profile, and passes the vibrations to the plate. Although we can have resonances between the ball mass and the equivalent contact spring, as described in the first rolling type, this does not dominate the sound. The resulting sound has thus the least temporal variation.




FIGURE 6.14: Slide and plate.

6.5 Comparison between measured and simulated plate acceleration

The experimental set-up used for the analysis of rebounds is now used again for comparing experiments with simulations of rolling spheres. The selected aluminum plate has a large waviness so that the other causes of roughness can be neglected, to a first approximation. All data of spheres and plate can be found in Table 6.2. These parameters were directly measured (m , V_0 , R , h , λ ,...), or extracted from the literature (Young's moduli, Poisson's ratios, thermoelastic constants). The Hertzian constant κ of the sphere-plate contact is estimated from the formula [58]:

$$\kappa = \frac{4\sqrt{R}}{3 \left[\frac{1-\nu_p^2}{E_p} + \frac{1-\nu_s^2}{E_s} \right]}, \quad (6.17)$$

where E_p, ν_p are the elastic parameters of the plate, and E_s, ν_s the elastic parameters of the sphere.

In Table 6.2, we can see that both spheres s2 and s4 have almost the same mass m , but that they differ in κ and R . Conversely, s2 and s5 have same the radius R , but differ in κ and m . Finally, s4 and s5 have similar values for κ , but differ in R and m . This selection was made intentionally with the goal to better understand the respective influence of size, mass and rigidity in the resulting sounds and vibrations. The initial velocity V_0 of the spheres is controlled by putting the sphere at height H_1 on a slide situated at one end of the plate (see Figure 6.14).

A rough estimate of the horizontal velocity at the beginning of the plate is then obtained by means of the simple formula:

$$V_0 = \sqrt{2gH_1}. \quad (6.18)$$

TABLE 6.2: Measured parameters used as input for the synthesis of rolling sounds in Section 6.5

Ball parameters:

Sphere	material	V_0	κ	R	m	E	ν
		m/s	$\times 10^9 \text{ N/m}^{2/3}$	mm	g	GPa	
s2	steel	0.53	7.8	10	32.6	200	0.3
s4	plastic	0.38	0.35	17	30.6	2	0.2
s5	plastic	0.64	0.27	10	3.5	2	0.2

Plate parameters:

Material : Aluminum ;

Density $\rho = 2.66 \text{ kg/m}^3$; Young's modulus $E = 63 \text{ GPa}$; Poisson's ratio $\nu = 0.35$;

Free boundary conditions ; mean thickness $\langle h \rangle = 8 \text{ mm}$;

Waviness amplitude $H_0 = 2 \text{ mm}$; waviness wavelength $\lambda = 5 \text{ cm}$;

Width $l_y = 20 \text{ cm}$; length $l_x = 100 \text{ cm}$;

Thermoelastic constants [12] $R_{th} = 8.45 \cdot 10^{-3}$ and $c_{th} = 8.0 \cdot 10^{-4} \text{ rad.m}^2.\text{s}^{-1}$.

Coefficients of the Padé approximant in the radiation damping term in Equation (6.2) [12]:

$a_0 = 1.1669$; $a_1 = 1.6574$; $a_2 = 1.5528$; $a_3 = 1$;

$b_1 = 0.0620$; $b_2 = 0.5950$; $b_3 = 1.0272$.

Coordinates of the observation point (accelerometer):

$x = 20.5 \text{ cm}$; $y = 5.0 \text{ cm}$.

If necessary, this estimation is refined by means of measurements with the high-speed camera. Comparison between measurements and simulations are made on acceleration signals rather than on sound pressure. The main reason is that the experiments are simpler and allow better reproducibility. Informal tests with a scanning microphone near the plate have shown that the waveforms and spectral contents of both signals are very similar, at a given position of the plate. The signal delivered by the accelerometer (B&K 4507-C) is amplified and sampled at a rate of 12800 Hz.

The parameters listed in Table 6.2 are used as input parameters for the synthesis. The calculation of 1.6 s of sound at a sampling frequency of 192 kHz ($3.07 \cdot 10^5$ samples) takes roughly 100 min on a standard desktop computer. The program yields, in addition, the displacement of the plate at any requested position, the contact force history and the displacement of the center of gravity of the sphere. The simulated acceleration at the observation point is then obtained through double differentiation versus time of the plate displacement. The analysis and post-processing of both experimental and simulated signals are made in MATLAB®. To allow comparisons with measurements, the simulated waveforms are first decimated by a factor of 15 (5×3). Thus, spectral analysis is made for frequencies between 0 and 6.4 kHz.

While listening to rolling sounds, the ear is sensitive to both the temporal and spectral features of the signal. The temporal envelope give us cues, for example, on amplitude modulation, presence of rebounds and presence of irregularities in the time history. The overall spectral balance is linked to our judgment of “dull” or “bright” sound. Finally, the fine structure of the spectrum, and the distribution of spectral peaks is of high importance in the perception of tone color and pitch. Taking advantage of our calculations in the time-domain, the results are presented both in the time (Figure 6.15) and frequency domain (Figure 6.16). As an alternative compromise, Figure 6.17 shows a time-frequency plot where both aspects are presents.

In Figure 6.15, the temporal envelopes are obtained through calculation of the energy of the signal. The energy is then low-pass filtered. This figure shows the results obtained for spheres s4 and s5. Measurements are on the top, and simulations are on the bottom. The initial velocity for sphere s4 (left-hand side of the figure) is slightly below the critical velocity: this results in only small modula-




FIGURE 6.15: Temporal envelopes of measured and simulated plate acceleration. First row: measurements ; second row: simulations. Left column: sphere 4 ; right column: sphere 5. The magnitudes are normalized to unity.




FIGURE 6.16: Average spectra of measured and simulated plate accelerations. First row: measurements ; second row: simulations. Left column: sphere 4 ; right column: sphere 5. Spectra are normalized to 0 dB.

tions of the contact force with only one or two clearly visible peaks due to a loss of contact. This effect is well reproduced in the simulations, although the exact position of the high peaks is not fully comparable: with such highly nonlinear signals, it is well-known that small differences in the initial conditions can lead to substantial differences after some time. The main reason for the discrepancies here is to find in the absence of “perfect control” of the initial conditions when the ball just reaches the plate after the slide. The initial velocity for sphere s5 (right-hand side of the figure) is well above the critical velocity: this results in much more pronounced peaks in the temporal envelope that account well for the fact that a lot of almost regular rebounds are present here. This envelope is clearly in accordance with the auditory perception, when listening to the corresponding sound. The number, shape, and periodicity of the peaks are well reproduced in the simulations. The overall distribution of amplitudes is also similar. As for sphere s4, the main discrepancy here is that there is no perfect reproduction of the amplitude of each peak. In addition, the “double rebounds” in the simulations are less visible in the measurements.

Figure 6.16 show the average spectra corresponding to the same acceleration signals shown in Figure 6.15. These spectra were obtained by computing the average power spectral density of the complete signal with windows of 8192 points and overlap of 8000 points. The frequency resolution is about 1.5 Hz. Compared to s4, the spectrum of s5 has clearly more energy in the upper frequency range (typically for frequencies above 2 kHz). This is in accordance with our daily experience when listening to rolling objects that are light, small and fast compared to others that are heavy, big and slow. Here, the Hertzian constant, which is comparable for both spheres, does not seem to play any role.

It can be seen that the simulations reproduce both the spectral envelope and the fine structure of the spectrum with great accuracy. All four figures show some “holes” in the spectrum around 2500 and 4500 Hz which is presumably caused by the position of the accelerometer. In the fine structure, we can recognize very narrow peaks, especially in the low-frequency range. From the recordings of the plate vibrations after the ball has left the plate, we could verify that these low-frequency “ringing” peaks correspond to excited eigenfrequencies of the plate. For the case of sphere s5, the plate is excited by rebounds, and we can see more peaks in the medium and upper frequency range, both in the measured and in the simulated acceleration. Despite its simplicity, we can see that our time-domain model is able to reproduce quantitatively, and with a relatively high level of detail, the fine structure of rolling sound spectra. Notice that the range of magnitude (in dB) is smaller for experimental signals compared to simulated ones, a consequence of the noise floor that is present in the measurements and not in the simulations.

The spectrograms shown in Figure 6.17 summarize the most significant audible features of a sphere rolling over a sinusoidal profile. Comparisons are made here between measurements and simulations for sphere s2. Both panels show similar patterns. On the “spectral” point of view, we can see that the energy is concentrated in the same frequency domains. In addition, sharp peaks, materialized here by very thin horizontal lines, are clearly visible with the same frequencies in both cases: as indicated above, these peaks correspond to eigenmodes of the plate. On the “temporal” point of view, we can see slow modulations marked here by large grey zones. On a finer time scale, the presence of superimposed rebounds is represented by small spots with the same duration both in simulations and experiments.




FIGURE 6.17: Spectrograms of the acceleration at the observation point while sphere 2 rolled over the plate. (Left) Spectrogram of the simulated acceleration ; (Right) Spectrogram of the measured acceleration.

6.6 Conclusion

In this chapter, a model for simulating the vibrations and sounds produced by spheres rolling on a damped elastic plate in the time-domain has been presented. This model includes a sphere-plate interaction based on Hertz's law of contact, a surface profile accounting for rugosity and/or waviness, and the flexural vibrations of the damped thin plate, using a Kichhoff-Love model.

One main objective of such simulations is to provide a tool for producing rolling sounds that can be heard and compared to the reality. Therefore, particular attention has been paid to the modeling of damping. For this reason, a previous model of single impact of a sphere against a plate, which has been proven to be efficient and yielding simulated sounds comparable to experiments, has been adapted to rolling sounds. On a numerical point of view, such an adaptation made it necessary to develop a new method of windowing for the interaction, in order to avoid discontinuities that would have led to unwanted noise in the simulations. In addition, the windowing associated to an interpolation procedure allows to deal with two different scales: a fine grid for the plate surface in order to compute the interaction force, and a larger grid for computing the flexural displacement of the plate. Such a procedure was necessary, since the use of the fine grid for both elements would have led to an overwhelming computation time.

The model has been validated by three successive procedures. First, simulations performed with a simple constant moving forces on a rail show the expected

Doppler effect associated to moving sources. A second procedure consisted in measuring the restitution coefficient of energy for balls bouncing on two different MDF plates and to compare the results with the simulations. Such a procedure was aimed at assessing that the global energy loss was correctly reproduced by the model and it gave satisfactory results. However, it was found that high sampling rates are necessary in order to obtain a very good agreement between measurements and simulations. In addition, some limits of the model appear for thick plates, for which the flexural model would probably need to be modified. Finally, comparisons were made between measurements and simulations of “critical velocities” for spheres rolling on an aluminum plate with a sinusoidal profile. The “critical velocity” corresponds to the threshold for which rebounds start to appear, leading to significant differences in the produced sounds. Here again, the model estimates such velocities correctly.

One main interest of such a physical model lies in its capability of providing a large variety of sounds, despite the fact that only 10 to 20 input parameters are necessary. To illustrate this, systematic simulations were conducted, in which only a few parameters are modified in each case, for a relatively large range of values. The four main classes of sounds obtained in this manner were presented in this paper. It has been proposed to designate them as: amplitude-modulated, periodic bouncing, chaotic bouncing and continuous contact.

Finally, comparisons were made between measured and simulated plate accelerations for different spheres rolling on a specially designed aluminum plate with a sinusoidal profile. The model is able to reproduce real temporal envelopes convincingly. The general distribution of amplitude and impulse duration is well reproduced, although the waveforms cannot be fully superimposed. To obtain perfect superposition, one would probably have to control the initial conditions much more carefully. Other causes of discrepancies might follow from microcontacts [21], or from a departure from straight line motion of the spheres. Spectral analysis of simulated and real sounds show an impressive number of similarities: here the spectral envelopes can be clearly superimposed, and even the detailed contents of the spectra match reasonably well. The auditory resemblance between real and simulated sounds is well rendered by spectrograms which show together the spectral formants, the ringing frequencies and the discontinuities in time.

In its present form, the model is theoretically able to account for discontinuities in thickness and inhomogeneities in the plate, but systematic exploration of such quantities remains to be done. Other challenges are linked to the objective of simulating “high-speed” rolling objects: in its present form, the model would then need to simulate very long plates and thus, there are certainly some alternative ways to find in order to limit the computation time. The contact model is also very elementary, and it is rather surprising that such a crude model yields so good results. However, one can anticipate that rolling of much more massive and large objects, such as wheels, would need a more sophisticated contact model, comparable to those developed in the context of wheel-rail interaction, including, for example, friction and plasticity [107] [51] .

6.A Numerical scheme

To solve the non-linear ball-plate contact problem, we use a fourth order Runge-Kutta-Nyström[2] approach, which we will describe here.

Let $f(W_p, W_b, \dot{W}_b)$ be the calculation of the acceleration using the mass of the sphere and Newton's second law, $\ddot{W}_s = F_s/m$ and Hertz' law in the case of contact between ball and plate.

Furthermore, $W_{ps}^n = W_{plate}^n + W_{surface}^n$ is the vertical position of the plate due to the vibration of the plate and the surface roughness, W_s is the position of the ball. Then the fourth order Runge-Kutta-Nyström equations read:

$$k1 = \frac{1}{2}\Delta t f(W_{ps}^n, W_b) \quad (A-1)$$

$$k2 = \frac{1}{2}\Delta t f(W_{ps}^{n+1/2}, W_s^n + K, \dot{W}_s^n + k1) \quad K = \frac{1}{2}\Delta t(\dot{W}_b^n + \frac{1}{2}k1) \quad (A-2)$$

$$k3 = \frac{1}{2}\Delta t f(W_{ps}^{n+1/2}, W_s^n + K, \dot{W}_s^n + k2) \quad (A-3)$$

$$k4 = \frac{1}{2}\Delta t f(W_{ps}^{n+1}, W_s^n + L, \dot{W}_s^n + k2) \quad L = \Delta t(\dot{W}_s^n + k3) \quad (A-4)$$

$$W_s^{n+1} = W_s^n + \Delta t(\dot{W}_s^n + \frac{1}{3}(k1 + k2 + k3)) \quad (A-5)$$

$$\dot{W}_b^{n+1} = \dot{W}_b^n + \frac{1}{3}(k1 + 2k2 + 2k3 + k4). \quad (A-6)$$

Note that in the case of no damping, $k2$ and $k3$ are equal. For the different plate positions, we take the same value, $W_{ps}^n = W_{ps}^{n+1/2} = W_{ps}^{n+1}$. Although W_{ps} changes due to plate vibrations and the changed position of the ball, taking this into account via a linear interpolation did not speed up the convergence of the restitution coefficient as a function of the sampling frequency as shown in Figure 6.10.

6.B Constant moving force

First we examine the effect of a movement of the source over the plate on the vibrations in the plate. This is best illustrated by a simulation that is otherwise as simple as possible, thus a propagating constant force without any modulations or interaction of the plate vibrations on the force. The simulated plate is very long, 11 m, to be able to view the wave propagation for a while before the first reflections occur. In the width direction, however, the plate is chosen to be small, 20 cm, to have a quasi-static situation. Thus we see the wave propagation only in one dimension. Five velocity profiles at 5.2 ms intervals are shown in Figure 6.18. The first panel shows the vibrational waves shortly after the start of the impact, which lies inbetween the low and highest peak in the middle of the figure. The source is moving leftward and we see that the plate moves downwards just before the impact point and upwards right behind the impact point. Here the vibrations are still very much centered around the impact point.

In Figure 6.18b we can see that the vibrational waves that are more outward, thus the ones that have traveled more distance, have shorter wavelengths than the ones more to the center. This effect, the higher frequencies having a higher wave velocity compared to the lower frequency waves, is known as dispersion. The Doppler effect can be identified the clearest in Figure 6.18c. The Doppler effect states that the waves in the direction of the source movement are spatially compressed (left in the figure) and the waves at the other side of the source are elongated (right in our figure). Indeed we can identify this difference in the pictures.

Due to dispersion, the lower the frequency, the lower the wave velocity. At a certain frequency, the propagation of the force source is faster than the propagation of the waves. The very slow waves stay approximately at the original point of impact, resulting in vibrations that are apparently steady on the plate. In fact, there exists exactly one frequency where the velocity of the vibrational waves is exactly equal to the speed of the propagating force. This results in a sinusoidal vertical movement of the impact point, despite of the fact that the force is constant.

Shortly after the last panel of Figure 6.18 the first reflected waves return to the center of the plate. With their increasing energy the picture gets more and more disturbed and we can no longer distinguish any of the previous effects.




FIGURE 6.18: Plate velocity due to a moving constant force on the plate as calculated by the numerical model. The plate is 11 m long and 20 cm wide, but only the central four meters are shown. The vertical velocity is amplified 20000 times. Until $t=0$, the plate is in rest. Then, the force starts at about the middle of the figure and propagates to the left with a speed of 40 m/s.




FIGURE 6.19: Overview of the analysed system with a non-linear spring.

6.C A simple non-linear system

In this section a simple non-linear system will be analysed to gain insight into the effects that can occur due to the non-linear behaviour of the contact between the ball and the plate. Specifically we are interested in the generation of vibration frequencies that are not present in the excitation and in the conditions in which loss of contact occurs.

Many simplifications will be made in this system. A one-dimensional Hertzian approximation was taken for the contact spring, while in reality the contact can be at more than one point [35]. The surface roughness is simply sinusoidal and the plate is considered completely rigid, which means also that no waves can travel through this plate.

6.C.1 A simple non-linear system with linear damping

We will try to explain some of the behavior of the rolling ball by analyzing this simplified non-linear system. The ball rolling over a plate is represented by a mass, and the plate has a rigid sinusoidally curved surface. The contact is modeled, according to Hertzian theory [46], with a non-linear spring.

Next, the conditions are identified for which the sphere loses contact with the plate. Similar systems were also studied using the method of the harmonic balance by Nayak [72], and by the method of multiple scales as used here by Hess and Soom [38] and by Turner [104]. Related experimental work was published by Rigaud and Perret-Liaudet [83].

The origin of the coordinates is taken at the contact point, before deformation, thus the ball stays at $x, y = 0$. The plate is moving with constant speed underneath the sphere, and, due to its surface roughness, there appears to be a variable displacement of the plate at the contact point. The system we analyze is a mass resting on a Hertzian spring, with an imposed displacement at the bottom of the spring. It is assumed that there remains contact between the sphere and the plate.

Hence $\forall t, z(t) < 0$. The plate impedance is taken to be infinite. As a first step, the forces on the system are listed, starting with F_m due to Newton's second law with ball mass M and vertical position z :

$$F_m = -M\ddot{z}, \quad (\text{A-7})$$

the gravitational force:

$$F_g = -Mg, \quad (\text{A-8})$$

and finally the non-linear spring according to Hertzian theory, but with a linear damping term added²:

$$F_s = \kappa(h - z)^{3/2} - 2M\mu(\dot{h} - \dot{z}). \quad (\text{A-9})$$

The surface profile with amplitude H_0 and wavelength Ω is given by

$$h = H_0 \cos(\Omega t). \quad (\text{A-10})$$

For the static case $z(t) = z_0$ we can neglect the damping term since $\dot{h} = 0$ and $\dot{z} = 0$, and the remaining terms, $F_s + F_g = 0$ leads to

$$\kappa(h - z_0)^{3/2} - Mg = 0, \quad (\text{A-11})$$

introducing the earlier used quantity, $\alpha = h - z$, as a positive

$$\alpha_0 = \left(\frac{Mg}{\kappa} \right)^{2/3}. \quad (\text{A-12})$$

For the dynamic case we add the driving position due to the surface, h , which indirectly results in a driving force via $F_s + F_g + F_m = 0$, leading to

$$\kappa(h - z)^{3/2} + 2M\mu(\dot{h} - \dot{z}) - Mg - M\ddot{z} = 0. \quad (\text{A-13})$$

Using the static result of Equation (A-11) then leads to

$$\kappa(h - z_0 + z_0 - z)^{3/2} - \kappa(h - z_0)^{3/2} - 2M\mu(\dot{h} - \dot{z}) - M\ddot{z} = 0, \quad (\text{A-14})$$

²The choice for a linear damping term seems the usual approach and evolves from practical considerations in the following calculations.

where the second, static, term $\kappa(h - z_0)^{3/2}$ does not cancel against terms of the first, dynamic, term $\kappa(h - z_0 + z_0 - z)^{3/2}$ as is the case for a linear spring. Using the Equation (A-12), we find

$$\kappa((\alpha_0 + z_0 - z)^{3/2} - \alpha_0^{3/2}) - 2M\mu\dot{\alpha} - M\ddot{z} = 0. \quad (\text{A-15})$$

Substituting $\bar{z} = z - z_0$, which is also equal to $z - h + \alpha_0$, thus gives $\dot{z} = \dot{\bar{z}} - \dot{h}$, and also $\dot{\bar{z}} = -\dot{\alpha}$, after which it results in

$$\kappa\alpha_0^{3/2} \left[\left(1 - \frac{\bar{z}}{\alpha_0} \right)^{3/2} - 1 \right] + 2M\mu\dot{\bar{z}} - M\ddot{\bar{z}} = -M\ddot{h}. \quad (\text{A-16})$$

If $\bar{z} \ll \alpha_0$ we can use the Taylor expansion,

$$\left(1 - \frac{\bar{z}}{\alpha_0} \right)^{3/2} = 1 - \frac{3}{2} \frac{\bar{z}}{\alpha_0} + \frac{3}{8} \left(\frac{\bar{z}}{\alpha_0} \right)^2 + \frac{1}{16} \left(\frac{\bar{z}}{\alpha_0} \right)^3 + \text{h.o.t.}, \quad (\text{A-17})$$

or, alternatively a third-order least-squares fit on the interval $0 < \bar{z}/\alpha_0 < 1$, similar to the fit used by Hess and Soom [38]. The advantage of this approximation is that the error is spread over the whole region, whereas the Taylor approximation is more precise for smaller values of $\bar{z}/\alpha_0$. The least squares fit can thus be expected to be more precise in the determination of loss of contact between the ball and the plate. Using the least squares fit we find

$$\left(1 - \frac{\bar{z}}{\alpha_0} \right)^{3/2} = \frac{384}{385} - \frac{16}{11} \frac{\bar{z}}{\alpha_0} + \frac{16}{77} \left(\frac{\bar{z}}{\alpha_0} \right)^2 + \frac{8}{33} \left(\frac{\bar{z}}{\alpha_0} \right)^3, \quad (\text{A-18})$$

however, we need the first term to be one to cancel out the other one within the square brackets of Equation (A-16). Therefore we make a small error and set the first term to one. The numeric constants of Equations (A-17) or (A-18) can be replaced by symbols c'_1 , c'_2 and c'_3 to obtain

$$\left(1 - \frac{\bar{z}}{\alpha_0} \right)^{3/2} = 1 - c'_1 \frac{\bar{z}}{\alpha_0} - c'_2 \left(\frac{\bar{z}}{\alpha_0} \right)^2 - c'_3 \left(\frac{\bar{z}}{\alpha_0} \right)^3. \quad (\text{A-19})$$

Combining this with Equation (A-16) results in

$$M\ddot{\bar{z}} - \kappa\alpha_0^{3/2} \left[-c'_1 \frac{\bar{z}}{\alpha_0} - c'_2 \left(\frac{\bar{z}}{\alpha_0} \right)^2 - c'_3 \left(\frac{\bar{z}}{\alpha_0} \right)^3 \right] + 2M\mu\dot{\bar{z}} = -M\ddot{h} \quad (\text{A-20})$$

$$\ddot{\bar{z}} + c'_1 \frac{\kappa}{M} \alpha_0^{1/2} \bar{z} + c'_2 \frac{\kappa}{M} \alpha_0^{-1/2} \bar{z}^2 + c'_3 \frac{\kappa}{M} \alpha_0^{-3/2} \bar{z}^3 + 2M\mu\dot{\bar{z}} = -\ddot{h}, \quad (\text{A-21})$$

or in the form used by Nayfeh and Mook [73, equation 4.2.1]

$$\ddot{z} + 2\mu\dot{z} + \omega_0^2 z + c_2 \bar{z}^2 + c_3 \bar{z}^3 = E(t) \quad (\text{A-22})$$

with

$$\omega_0^2 = c'_1 \frac{K}{M} \alpha_0^{1/2} \quad (\text{A-23})$$

$$c_2 = c'_2 \frac{K}{M} \alpha_0^{-1/2} \quad (\text{A-24})$$

$$c_3 = c'_3 \frac{K}{M} \alpha_0^{-3/2} \quad (\text{A-25})$$

$$E(t) = H_0 \Omega^2 \cos \Omega t. \quad (\text{A-26})$$

Here we can identify a shortcoming of the least squares fit method. The ω_0 can be calculated easily by linearizing the Hertz contact spring around the static equilibrium as shown by Sabot *et al.* [86]. The analytical solution obtained in this way is exactly equal to the ω_0 found by the Taylor approximation, but the least squares method results in a value 3 percent lower.

To solve this equation using the method of multiple scales, the problem is ordered in different scales, using a dimensionless parameter, thus replacing $\bar{z}$ by ϵu , and in order to obtain the forcing and damping terms in the smallest scale, because the damping has only a small influence, they are multiplied by ϵ^3 . Dividing all by a common single ϵ we obtain:

$$\ddot{u} + \omega_0^2 u = -2\epsilon^2 \mu \dot{u} - \epsilon c_2 u^2 - \epsilon^2 c_3 u^3 + \epsilon^2 k \cos \Omega t. \quad (\text{A-27})$$

Now the single function u can be replaced by a sum of different functions, one for each of the time scales $T_n = \epsilon^n t$, with $n \in \mathbb{N}$,

$$u(t, \epsilon) = u_0(T_0, T_1, T_2, \dots) + \epsilon u_1(T_0, T_1, T_2, \dots) + \epsilon^2 u_2(T_0, T_1, T_2, \dots) + \dots \quad (\text{A-28})$$

For the differential operators, we find

$$\frac{d}{dt} = \frac{dT_0}{dt} \frac{d}{dT_0} + \frac{dT_1}{dt} \frac{d}{dT_1} + \frac{dT_2}{dt} \frac{d}{dT_2} + \dots = D_0 + \epsilon D_1 + \epsilon^2 D_2 + \dots, \quad (\text{A-29})$$

and for the second-order differential operator all terms with ϵ of order higher than 2 can be excluded, the remainder being

$$\frac{d^2}{dt^2} = D_0^2 + 2\epsilon D_0 D_1 + \epsilon^2 (D_1^2 + 2D_0 D_2). \quad (\text{A-30})$$

Finally, we assume that the driving frequency is close to the free vibration frequency, by replacing Ω by $\omega_0 + \sigma\epsilon^2$. Combining all and substituting the result into Equation (A-27), and separating the different scales gives

$$D_0^2 u_0 + \omega_0^2 u_0 = 0 \quad (\text{A-31})$$

$$D_0^2 u_1 + \omega_0^2 u_1 = -2D_0 D_1 u_0 - c_2 u_0^2 \quad (\text{A-32})$$

$$D_0^2 u_2 + \omega_0^2 u_2 = -2D_0 D_1 u_1 - 2D_0 D_2 u_0 - D_1^2 u_0 - 2\mu D_0 u_0 - 2c_2 u_1 u_0 - c_3 u_0^3 + k \cos(\omega_0 T_0 + \sigma T_2). \quad (\text{A-33})$$

These three formulas at the three different scales will now be solved one after another. For the first equation, the solution reads

$$u_0 = A \exp(i\omega_0 T_0) + \bar{A} \exp(-i\omega_0 T_0), \quad (\text{A-34})$$

where A is an unknown complex function with $A = A(T_1, T_2)$ and $\bar{A}$ the complex conjugate of A . The governing equations for A are obtained by requiring u_1 and u_2 to be periodic in T_0 . The solution for u_0 can now be included in Equation (A-32) to obtain

$$D_0^2 u_1 + \omega_0^2 u_1 = -2D_1 A i\omega_0 \exp(i\omega_0 T_0) - c_2 (A^2 \exp(2i\omega_0 T_0) + A\bar{A}) + cc, \quad (\text{A-35})$$

where cc indicates the complex conjugate of the preceding terms. In order to cancel the first part, which is secular, we obtain $D_1 A = 0$ or $A = A(T_2)$. This results in

$$u_1 = \frac{c_2}{\omega_0^2} [-A\bar{A} + \frac{1}{3} A^2 \exp(2i\omega_0 T_0) + \frac{1}{3} \bar{A}^2 \exp(-2i\omega_0 T_0)]. \quad (\text{A-36})$$

Substituting these results for u_0 and u_1 into Equation (A-33), and dropping all terms containing $D_1 A$ and those periodic in $3\omega_0 T_0$ results in

$$D_0^2 u_2 + \omega_0^2 u_2 = \left[-2\mu i\omega_0 A - 2\frac{c_2^2}{\omega_0^2} A^2 \bar{A} - \frac{4c_2^2}{3\omega_0^2} A^2 \bar{A} - 3c_3 A^2 \bar{A} + \frac{1}{2} k \exp(i\sigma T_2) - 2D_2 i\omega_0 A \right] \exp(i\omega_0 T_0) + cc + NST, \quad (\text{A-37})$$

where NST stands for non secular terms. In this case we are only interested in the conditions to cancel the secular terms. Thus,

$$2i\omega_0 (\mu A + D_2 A) + \frac{10}{3} \frac{c_2^2}{\omega_0} A^2 \bar{A} - 3c_3 A^2 \bar{A} - \frac{1}{2} k \exp(\sigma T_2) = 0. \quad (\text{A-38})$$

To find these conditions, we rewrite A in polar form, $A = \frac{1}{2}a \exp i\beta$ where a and β are real functions of T_2 , therefore $D_2A = \frac{1}{2}a' \exp(i\beta) + \frac{1}{2}ia\beta' \exp(i\beta)$ resulting in

$$i\omega_0(\mu a \exp(i\beta) + a' \exp(i\beta) + ia\beta' \exp(i\beta)) + \frac{1}{2^3} \left(\frac{10}{3} \frac{c_2^2}{\omega_0} - 3c_3 \right) a^3 i \exp(i\beta) - \frac{1}{2} k \exp(i\omega_0 T_0) = 0, \quad (\text{A-39})$$

separating real and imaginary terms,

$$\begin{cases} a' &= -\mu a - \frac{k}{2\omega_0} \sin(\sigma T_2 - \beta) \\ a\beta' &= \frac{10c_2^2 - 9c_3\omega_0^2}{24\omega_0^3} a^3 - \frac{k}{2\omega_0} \cos(\sigma T_2 - \beta) \end{cases}, \quad (\text{A-40})$$

which results in the general approximate solution

$$u = a \cos(\Omega t - \gamma) - \frac{1}{2} \epsilon \frac{c_2}{\omega_0} a^2 \left[-1 + \frac{1}{3} (2\Omega t - 2\gamma) \right] + O(\epsilon^2), \quad (\text{A-41})$$

where $\gamma = \sigma T_2 - \beta$. The equation can represent both softening and hardening behavior, but the condition for hardening behavior is never fulfilled. Hardening behavior implies, $c_3 > 10/9 c_2^2 \omega_0^{-2}$, which using Equations (A-23)-(A-25), equals

$$c'_3 > \frac{10}{9} \frac{c_2'^2}{c_1'}, \quad (\text{A-42})$$

which is in our case is never fulfilled since c'_3 is negative and c'_1 is positive. For the lift of the mass, $-\frac{1}{2} \epsilon c_2 \omega_0^{-2} a^2$, we get

$$\text{lift} = -\frac{1}{2} \frac{\kappa}{M \omega_0^2 \sqrt{\alpha_0}} c'_2 a^2. \quad (\text{A-43})$$

The ϵ was removed from the formula because it was only introduced to order the different scales of the formula. The minus sign indicates that the deformation is reduced, and the mass is moved upwards. For conditions that lead to loss of contact, the amplitude a must be near the static deformation α_0 . In fact, the loss of contact occurs a little earlier due to this lift, but if the lift is small compared to the oscillation amplitude we can substitute a by α_0 and obtain a shorter expression

$$\text{lift} = -\frac{g c'_2}{2 \omega_0^2}. \quad (\text{A-44})$$

Next we look at the detuning parameter, σ . This parameter indicates the change in resonance frequency for changing amplitudes, which is typical for non-linear systems. No explicit solution is given in by Nayfeh and Mook but the derivation is very similar to one of the simpler cases, thus, we take from the solution

$$a' = -\mu a + \frac{k}{2\omega_0} \sin \gamma, \quad (\text{A-45})$$

$$a\gamma' = a\sigma - \frac{9c_3\omega_0^2 - 10c_2^2}{24\omega_0^3} a^3 + \frac{k}{2\omega_0} \cos \gamma. \quad (\text{A-46})$$

The primed terms can be set to zero to find the stationary solution. We will solve σ from these equations. Starting by bringing the non-circular terms to the left then squaring and summing the last two equations results in

$$\left[\mu^2 + \left(\sigma - \frac{9c_3\omega_0^2 - 10c_2^2}{24\omega_0^3} a^2 \right)^2 \right] a^2 = \frac{k^2}{4\omega_0^2}, \quad (\text{A-47})$$

$$\sigma = \frac{9c_3\omega_0^2 - 10c_2^2}{24\omega_0^3} a^2 \pm \sqrt{\frac{k^2}{4\omega_0^2 a^2} - \mu^2}, \quad (\text{A-48})$$

which is the frequency-response curve. The first part will give the free vibration spline or backbone. The backbone represents the resonance frequency as a function of oscillation frequency in the absence of damping. With the use of our c_1 and c_2 , the backbone is

$$\sigma = \frac{27}{32 \cdot 24} \left(\frac{3}{2} \right)^{3/2} \frac{\kappa^{13/12}}{M^{13/12} g^{7/12}} a^2. \quad (\text{A-49})$$

The backbone and frequency response curve are plotted in Figure 6.20. We can see the softening behavior, i.e., the backbone is curved towards lower frequencies for higher amplitudes. By comparing the oscillation amplitude with the static deformation, the plotted horizontal dashed line in Figure 6.20, we can see if loss of contact occurs. In the example shown this is not the case. When the damping μ is lower, or the excitation is higher, the frequency response curve will become higher than the amplitude of loss of contact, and loss of contact will occur.

6.C.2 Conclusion

This analytical calculation regards a very simple setup that has many simplifications. We can view upon this calculation as a generalization of Falcon's result,




FIGURE 6.20: The frequency response curve of the non-linear system. The solid line gives the vibration amplitude for the ball given the excitation amplitude of 1 mm. The dotted line represents the backbone. The dashed line represents loss of contact. For the used $\mu = 1 \text{ s}^{-1}$ and $H_0 = 1 \text{ mm}$ there will be no loss of contact. When damping is lower or the driving amplitude is higher, however, loss of contact will be possible. The oscillation frequency calculated by Falcon *et al.* [20] is indicated with the circle. This point should be equal to the crossing of the backbone curve (dotted) and the loss of contact (dashed).

who calculated the oscillation frequency near loss of contact [20] using an energy balance method.

This calculation can aid to obtain more understanding about the nature of the ball plate contact. At least for this simple plate we can calculate the speed for loss of contact. This is not a practical result, but the oscillation frequency was used to verify the simulation program in the preceding chapter. In the inverse situation, when a ball comes to rest after having bounced on a plate we also find the same oscillation frequency [20].

7 | Using Psychometric Functions to Determine the Auditory Capability to Perceive the Size and Speed of a Rolling Ball¹

Abstract

The purpose of the experiments presented in this chapter was to quantify the precision of listeners' judgments of certain parameters of a mechanical system from the sounds generated by this system. We used sounds generated by a wooden ball rolling over a wooden plate. In the recordings of these sounds, the ball diameter, the rolling speed and the thickness of the plate were varied. Two experimental methods were used: In the paired-comparison experiment, participants listened to two sequentially-presented sounds and had to indicate, which of the two sounds was, e.g., created by the larger ball. Corresponding measurements were performed for paired comparisons of the mechanical parameters ball speed and plate thickness. In the absolute magnitude-estimation experiment, participants listened to single sounds and had to estimate the value of, e.g., the size of the ball by reporting a number. Again, also the two other parameters were tested. The data obtained with both methods were used to calculate psychometric functions, which indicate the listeners' sensitivity to differences in one of the physical parameters. For each pair of sounds the percentage of correct responses was plotted against the logarithm of the proportion of their mechanical parameter values. These datapoints were then fitted by a cumulative normal distribution. Both paired comparison and absolute magnitude estimation led essentially to the same psychometric functions for the size estimations. In the psychometric functions for the speed estimation, however, there was a difference between the results obtained with the two methods, and the origin of this discrepancy will be discussed. By applying this analysis to subsets of the data it is shown that the

¹Part of this work was presented at the ASA Providence meeting in 2006 [95]

precision of the listeners' estimations of the ball size changes when the ball speed changes. The same is true for the estimations of the ball speed when the ball size changes. In contrast to the results for judgments of the two *ball* parameters, the plate thickness was poorly judged by the participants in the paired comparison task.

7.1 Introduction

The vibrations produced by a mechanical system convey information about its physical properties, which are passed to the listener via sound. Our main goal is to estimate how much and what type of information is passed to the listener via these sounds. We want to study, therefore, how precisely the listener can judge certain mechanical properties of the system by listening to the sounds, how much difference there is between participants, and whether listeners can hear and independently judge more than one property of the system at a time.

We used sounds generated by a wooden ball rolling over a wooden plate. These continuous sounds present listeners with different kinds of information at the same time: The size of the ball, speed of the ball, and the plate thickness, which were all varied during the experiments.

Participants were asked to judge each of these three parameters in the sound in two different ways, in a two-alternative forced-choice paired-comparison task, and using absolute magnitude estimation. This method was used to test whether discrimination capabilities decrease when two parameters, e.g., speed and size, are changed simultaneously.

The auditorily perceived speed of objects passing by was studied by Kaczmarek [49]. He discussed three types of cues for determining the speed of these objects, namely, interaural time differences, loudness variations and the Doppler shift. The first is not present in our situation, because we presented the sounds monaurally to the participants. Another difference with our experiments is that in our case it was the plate that radiated the sound and not the moving object, the ball. This is due to the fact that the plate is much larger and radiates more efficiently. The plate, however, did not move. When the microphone is placed too close to the plate, the recorded sound will be the loudest when the ball rolls close to the microphone. We choose our microphone placement so that this effect was avoided.

Much research has been done on the subject of bouncing balls and impacted plates. Examples of the auditory perception of mechanical properties of a system

with a ball, or more general an impactor, and plate are: The complex thickness patterns [16], the shape of plates [54], and the size of bouncing balls [31] [96].

Psychometric functions have been used to determine the capabilities of listeners to distinguish differences in the frequency spectrum of the sound [64]. We will apply this method in the field of psychomechanics to determine the capabilities of listeners to distinguish mechanical differences in the sound generating system, in our case the speed and size of rolling balls. In this respect it is a generalization of the result of Houben *et al.* [39][42][43], who used a paired-comparison method to study the perception of the size and speed of rolling balls.

We will compare two methods for gathering the data, namely paired-comparison and absolute magnitude estimation. The results were sometimes similar and sometimes different. We will propose an explanation for these discrepancies based on different processing levels of the acoustic information.

7.1.1 Stimuli

The stimuli consisted, as mentioned, of recordings of wooden balls rolling over wooden plates. In the recording of these sounds, the ball diameter, the rolling speed and the thickness of the plate were varied independently. The balls and plates were chosen to give good rolling sounds and to be the same or similar as used in earlier experiments [40].

The balls we used were made of beech wood and had diameters of 25, 35, 45, 55 and 68 mm. Due to the production process, the balls were not perfectly spherical, and this leads to temporal variations in the sound made as they roll, such as amplitude modulation or even some bouncing for the smallest and fastest balls.

In order to give the balls a well defined velocity along the center line of the plate, they were rolled from a slide onto the plate. The height of the slide was 25 cm. The balls were released from various heights on the slide to vary their velocity. In order to measure the velocity of the ball, its course interrupted the beams of six independent light-gates, placed at intervals of 20 cm. The measurements showed that, during one run, the ball's velocity was constant within a few percent. The velocities we used in the experiment ranged from 0.48 to 1.63 m/s.

The slide was separated from the plate by a narrow slit, about 0.5 mm in width, to avoid the transmission of vibrations from the slide to the plate. The slide had a smooth bend near its end to transform the vertical velocity of the balls into a horizontal one, and so minimizing the impact heard when the ball hits the plate. At the other side of the plate, the ball was free to roll off.

The plates, 50 cm wide and 120 cm long, were made of MDF, which consists of wood-particles pressed and glued together. Their thickness was 4, 8, 12 or 18 mm. The advantage of this kind of material over real wood is that the material properties do not change with position on the plate.

Seen from above, the microphone was placed 70 cm to the right of the long side of the plate, and halfway along its length. It was placed 70 cm above the plate, pointing at it with an angle of 45 degrees. The closest distance between microphone and plate was 100 cm. The microphone itself was a small diaphragm cardioid condenser microphone, type Røde NT5. It was connected to a sound card, type Terratec MIC8, that recorded the sounds with 24-bit resolution at a sample frequency of 48 kHz on a standard PC.

From the middle of the sound, 600 ms segments were cut out. The waveforms were gated on and off with half of an 8 ms Hanning window. In the paired comparison task of experiment 1, there was 250 ms of silence between the two sounds. The sounds were presented to the participants via headphones, at their original relative sound levels.

Because spectral information is often considered to be one of the most important cues for determining object properties, a plot of the average spectrum of some the used sounds, with different ball speeds, ball sizes, and plate thicknesses is included in Figures 7.1 to 7.3. The method used to obtain these spectra is Welch's averages periodograms [79] with a window length of 512 samples or 10.7 ms. The spectrum of one sound is depicted in all three figures, a ball with a diameter of 45 mm, rolling with 0.91 m/s on a 12 mm thick plate, but in each figure two, or three for the plate thickness, other sounds with changes in one dimension are added: A smaller and larger ball in Figure 7.1, a faster and slower ball in Figure 7.2, and a ball rolling over a thicker and two balls rolling over a thinner plate in Figure 7.3. We can see that size and speed changes have somewhat similar effects on the spectrum, namely the amplitude of high frequencies increases and that of low frequencies diminishes for smaller and also for faster balls. The plate thickness has a different effect on the spectrum. The low frequencies are less pronounced for the thickest plate and the fine structure changes due to different modal frequencies of the plate. Although this cannot be seen indisputably from these figures, from general acoustics, see for instance [48], we know that the resonance frequencies of the plate change with its thickness, whereas the excitation only changes the amplitude of the vibrational modes.




FIGURE 7.1: Effect of the size of the ball on the average spectrum of the rolling ball sound. In all three cases the ball speed was 0.91 m/s and the plate thickness was 12 mm.




FIGURE 7.2: Effect of the speed of the ball on the average spectrum of the rolling ball sound. In all three cases, the ball diameter was 45 mm and the plate thickness was 12 mm.




FIGURE 7.3: Effect of the plate thickness on the average spectrum of the rolling ball sound. In all three cases, the ball speed was 0.91 m/s and the ball diameter was 45 mm.

7.2 Experiment 1: Paired comparisons

Experiment 1 was similar to earlier experiments described by Houben *et al.* [42]. This experiment was repeated in order to realize that we can compare the results of this experiment with those of experiment 2. Some extra pairs were added to spread the measurement points on the psychometric functions more evenly. Another difference was that we asked participants to judge the distance the ball traveled during the course of rolling, instead of asking its speed directly. When asked for speed, participants can choose to interpret this as "angular velocity" or as "linear velocity" [42], which gives different results. For brevity, and because this is really the underlying property that we are interested in when we ask for the traveled distance, we will still refer to the speed of the ball in the analysis. The actual speed can simply be calculated from the distance by dividing the distance by the sound sample duration, in this experiment 600 ms.

7.2.1 Method

The experimental methodology was a two-alternative forced-choice procedure. Thus the sound pair was presented once and the participants were asked to judge which of the two sounds represented the largest ball, which ball covered the largest distance, or which ball rolled over the thickest plate. They were then allowed to answer if they thought it was the first or second sound of the pair.

Feedback was never given. First, six practice sounds were presented to the listeners to get them accustomed to the procedure. At three points during the experiment, the participants were allowed to take a break. The experiment took about 25 minutes and the 23 participants were paid 3.50 euros for participating. They were recruited at the university and were mainly students. All participants reported that they had no hearing problems.

The first experiment consisted of three parts. In the first part the participants were asked to compare the diameters of the two balls and indicate the larger one. All combinations of the five different ball sizes were used, giving $5 \times 4 = 20$ pairs of sounds, including all pairs in both orders but not those consisting of two equal sounds. All four different ball speeds were used with one plate thickness, 12 mm. The two sounds always differed only in ball size. This results in a total of 80 sound pairs for the first part.

In the second part, the participants were asked to compare the distances traveled by the two balls and indicate the longer one. All combinations of the four different ball speeds were used, giving $4 \times 3 = 12$ pairs of sounds. All five different ball sizes were used with one plate thickness, 12 mm. The 60 sound pairs in which two sounds differed only in ball speed were followed by 24 pairs, consisting of the same ball speeds, but in which the first ball had a diameter of 25 mm, the second 68 mm and vice versa. This gives a total of $60 + 24 = 84$ sound pairs for the second part of the experiment.

In the third part of the experiment, the participants were asked to compare the thicknesses of the two plates on which the balls rolled and indicate the thicker one. All combinations of the four different plate thicknesses were used, giving $4 \times 3 = 12$ pairs of sounds. All five different ball sizes were used and one ball speed, 0.91 m/s. The two sounds of one pair always differed only in plate thickness. This results in a total of 60 sounds pairs, which were presented twice to the participants, for the third part. Together with some sound pairs of which the results are not used in this paper, the participants were asked to judge 372 sound pairs.

7.2.2 Results

To check if participants were able to perform the task, we plotted the average percentage correct for each type of sound pair, and for each individual participant. Figures 7.4 to 7.6 show these results for each of the three object properties that the participants were asked to judge: The size and speed of the ball and the thickness

of the plate it rolls over. In addition, the 95% confidence interval for guessing is shown between dashed lines in the figure.²

In Figure 7.4 the percentages correct for the size judgment are shown. It can be seen that all participants were able to do this task, although some were much better at it than others. One participant, number 14, performed very well in the task, with nearly 100% correct score. Overall, the average percentage correct for the size judgment was 85%. Since all participants scored above chance level, and on average scored well, we can conclude that the participants were able to do this task. In Figure 7.5, similar performance can be seen for the speed judgments. However, three of the participants performed well below chance level, i.e., they gave the wrong answer most of the time. These are called inverse-labelers, since they are apparently capable of identifying the speed difference, but they give the wrong label, i.e., they label the faster ball as slower. It can also be seen that two of these inverse labelers, participants 9 and 10, are among the best performers of all participants, giving consistent answers in 90% of the time. The results of these inverse-labelers were therefore inverted and included in the remainder of the analysis. We could also have chosen to leave out these three inverse labeling participants, but this leads to about the same results in any of the analyses, simply because these three participants together are, by chance, average responders in our tests. This taken into account, the average percentage correct for the speed judgment was 83%. The tested ranges for speed and size are more or less comparable; the ratio of the diameter of the largest ball and the diameter of the smallest one is 2.7 for the speed judgment. The ratio of the largest and smallest traveled distance was 3.4. From these figures it might seem that the discrimination of size and speed is about equal. This, however, is not the case as we will see later. The fact that the average percentage correct for the size and speed judgments are so close (85% and 83% respectively) is because, in the design of the experiment, the presented sound pairs contain, in both experiments, about an equal number of easy and difficult to distinguish stimulus pairs.

The analysis of the thickness judgment data, shown in Figure 7.6, however, shows a different picture. As many as ten participants fall within the confidence intervals for guessing, four of them fall below the confidence interval for guessing and only nine of them show behavior in line with good discrimination and correct labeling. None of them is outstanding. The four participants who score

²If a particular participant was not able to do the task for at least some of the sound pairs, it is assumed that the participant has guessed, and thus the 95% confidence intervals can be calculated by taking the 2.5% and 97.5% value of a binomial distribution for a process with probability parameter $p=0.5$ and the given number of soundpairs in the set.




FIGURE 7.4: Percentage correct of the answers of the 23 participants who were asked to indicate the larger ball in the pair of two presented sounds of rolling balls. Both balls had the same speed, four different speeds were used, and rolled over the plate with 12 mm thickness. The horizontal dashed line represents the upper and lower boundary of the 95% confidence interval for guessing. All participants fall outside this region.

below the confidence interval thus give more or less consistently the wrong answer. Their scores, however, are not set apart from the rest as was done case in the speed judgments. We could treat the data from these inverse-labelers in various ways; we could include their data as-is, inverse their answers or exclude their data. If we combine the data of all participants, however, performance differs significantly from guessing, regardless of how we treat the inverse-labelers. We can conclude, therefore, that, on average, the participants performed very poorly, but were still capable of judging above chance the thickness of the plate. Although it was possible to calculate a psychometric function from these data, one may then ask which participants to include, e.g., must the inverse labelers be included, or should we include only the participants who score above the 95% confidence limits. The thickness judgment data are, therefore, not analyzed in the remainder of this chapter.

7.2.3 Psychometric functions

We can estimate the psychometric function from the percentage correct data by plotting the percentages correct, as shown in Table 7.1 against the difference between the two sounds on one of the measured variables. It turns out that this




FIGURE 7.5: Percentage correct of the answers of the 23 participants who, after listening to the sounds of two rolling balls, were asked to indicate the ball which traveled the greatest distance in the course of rolling. The two balls had both the same size, out of five different sizes used in the experiment, and rolled over the plate with 12 mm thickness. The horizontal dashed lines represent the upper and lower boundary of the 95% confidence interval for guessing. All participants fall outside this region, but three of them reversed the fastest and slowest ball.




FIGURE 7.6: Percentage correct of the answers of the 23 participants who, after listening to the sounds of two rolling balls, were asked to indicate the ball which rolled over the thickest plate. The two rolling balls had both the same size and speed. The horizontal dashed lines represent the upper and lower boundary of the 95% confidence interval for guessing. Although the overall average percentage correct, 53.7%, is above chance level, 51.8%, participants clearly do not do well in this test.

TABLE 7.1: The percentage of answers: “ball 1 is larger” for all different stimulus pairs, averaged over all subjects. From this table the psychometric function shown in Figure 7.9 is derived.

68	0.93	0.95	0.83	0.79	–
55	0.96	0.88	0.67	–	0.13
45	0.91	0.88	–	0.3	0.08
35	0.92	–	0.23	0.15	0.08
25	–	0.42	0.16	0.08	0.03
	25	35	45	55	68

actual ball2 size [mm]

leads to the most consistent results if the logarithm of the proportion of the ball sizes is taken as the difference in physical variable. The fact that we need to use the difference of the logarithm of the ball sizes agrees with Steven’s law which predicts that the perceptual distance between stimuli differing in one physical dimension depends linearly on the difference of the logarithm of this quantity. Furthermore, if the datapoints are not plotted on a logarithmic scale, but on a linear scale the plots look no longer symmetric. On the other hand, the data points are symmetric on the x-axis if we define the distance between two sounds as the difference of the logarithm of the two ball sizes, as can be seen in Figures 7.7 and 7.9. Through the points found in this way, we fitted a cumulative Gaussian curve using a least-squares algorithm. The data were tested for response bias by checking the number of times the participants chose the first and the second sound as being the larger, the faster, or the ball rolling over the thicker plate. The data turned out to be symmetric with respect to the order of presentation and within chance levels for all tests that we ran. The psychometric curves can, therefore, be assumed to be symmetric around the point $(0, 1/2)$, for which the mechanical parameters in the two sounds are equal. Furthermore, we assumed that, for cases in which the physical differences of the balls are very large, i.e., larger than in the pairs we presented here, participants will always give the correct answer. This causes the psychometric curve to take values ranging from exactly zero to one

but not including zero and one. In this way, the cumulative Gaussian curve had only one degree of freedom. From the least-squares fit we can also find the confidence intervals for the estimated parameter, in our case β which is the inverse of the standard deviation of the fitted distribution, and represents the steepness of the psychometric function. For a comparison between different types of curves used for fitting psychometric functions, see [27], where our type of curve is listed as a Gaussian function with log-ratio kernel.

In contrast with the fitted psychometric functions, the data-points shown in the psychometric curves of the first experiment are not exactly symmetrical. As an example we take the psychometric function for size estimation presented in Figure 7.9. If the smallest ball was the first of the two sounds presented to the listeners, the result will be plotted in the left side of the plot, if it was presented second, this datapoint would appear in the right half. Since we found no dependence on the order of presentation, we could mirror the data to reduce the variance on the data. However, in the procedure used this would lead to exactly the same psychometric functions.

The results for the estimation of the speed are depicted in Figures 7.7 and 7.8. In Figure 7.7 we see the results for the speed judgment where the size of the two balls was equal. In Figure 7.8 the results are shown for speed judgments, in which the size of the balls in a pair also differed. We then see that the steepness of the psychometric function is much lower when the two balls also differ in size. The size judgment itself has the highest steepness, as shown in Figure 7.9.

7.3 Comparison with other studies

A two-alternative forced-choice experiment in a similar setup, but with different sound recordings, has been conducted by Houben [40]. In this section we will compare these data with ours. In Figure 7.10 we see a plot of the psychometric function for the estimated ball size, together with data from Houben. The data from the previous study of Houben fall somewhat around our data points. The somewhat larger variance can be explained by the fewer number of subjects in the earlier study, 6 versus 23 in our study. When comparing the speed data of both studies, the story is different. The data from Houben's experiment result in a 6.5 times greater steepness of the psychometric function, as shown in Figure 7.11. There are a few possible causes for this difference. First, the setup was different. Particularly the plate and plate support were different. In Houben's experimental setup the plate was more damped by the support, a layer of felt. Furthermore he




FIGURE 7.7: The psychometric function for the estimated differences in speed when the ball sizes were equal. The parameter β is the fitted steepness parameter, and the numbers between brackets present the 95% two-sided confidence intervals of the fit. Data are from the paired-comparison task.




FIGURE 7.8: The psychometric function for the estimated differences in speed when the ball sizes were different. The parameter β is the fitted steepness parameter, and the numbers between brackets present the 95% two-sided confidence intervals of the fit. Data are from the paired-comparison task.




FIGURE 7.9: The psychometric curve for the estimated differences in size when the ball speeds were equal. The parameter β is the fitted steepness parameter, and the numbers between brackets present the 95% two-sided confidence intervals of the fit. Data are from the paired-comparison task.




FIGURE 7.10: The psychometric function for the estimated size of the ball. The data from the current experiment, shown with filled dots, are compared with data from Houben [40], plotted with open dots.




FIGURE 7.11: The psychometric function for the estimated speed of the ball. The data from the current experiment, shown with filled dots, are compared with data Houben [40], plotted with open dots.

selected rolling sounds for which the ball was not bouncing too much. These two factors could lead to a more pronounced amplitude modulation in the sound of these rolling balls. Also, Houben minimized the bouncing by selecting the most spherical balls, and rolling directions for these balls, and by selecting recordings that, by chance, had the least bouncing in them. Sounds containing less bouncing may have been less confusing to the listeners. Although bouncing can provide information about the rolling speed, it could be difficult for the listener to compare a bouncing ball with a non-bouncing one, as may have been the case in our study.

Another possibility is the change of the question, posed to the participants, which was directly for speed in Houben's experiment and for traveled distance in the current one. The question was changed because the term "speed" could be interpreted in different ways. An example is given by Houben [40]. When we measure speed in m/s then a mouse that runs "fast", corresponds to a much lower speed than a horse that runs "fast". In other words, for a large object we expect a larger speed than we expect for a small object, to label it fast. For rolling balls it might be that listeners think a ball rolls fast, when it rotates fast instead of when it moves fast. But, this has the same effect, that larger balls should move faster than smaller balls to obtain the label fast. Consequently, the two effects are indistinguishable. The term rotational speed, the number of rotations of the ball per second, will therefore be used to indicate that a larger balls should move




FIGURE 7.12: The psychometric function for the estimated speed of the ball. The data from the current experiment, shown with filled dots, are compared with data from a pilot experiment, shown with open dots, by Dr. Laurie Heller, Dept of Cognitive and Linguistic Sciences, Brown University, Providence RI.

faster than smaller balls to obtain the label fast. To move at the same speed, measured in m/s, which will be called linear speed, a ball being twice as large as a reference ball should have half the rotational speed of the reference ball.

If listeners have no intuitive feeling for the traveled distance, the task could be much more difficult for the participants in the new experiment. They might have an intuitive idea about the rotational speed. They are, however, forced to answer in terms of linear speed when we ask for “traveled distance”. To come up with this answer, listeners will have to calculate the linear speed from the rotational speed using the perceived size of the ball. The perceived ball size is, however, not completely accurate and neither is the rotational speed. Hence, in the process of these calculations there are two sources of inaccuracies. Therefore, this process could have led to the lower performance we observed compared to the performance reported previously.

This, however, is not likely to be explanation for the much lower steepness of the psychometric function in our study, because we also found the same difference in the scales of an multidimensional scaling (MDS) experiment, as will be discussed in Chapter 8 of this thesis. As we explain in Appendix A of this chapter, the differences in the scale sizes of the MDS analysis could cause a difference in detectability. A numerical calculation leads to the observed difference in the steepness of the psychometric functions.

Further evidence that the question posed is not of crucial importance, comes from a pilot study by Heller at Brown university (personal communication). Four participants took part in a two-interval forced-choice test. This experiment was equal to ours in the sense that the sound samples were the same as in our experiments. The question posed to the participants was, however, different. We asked the participants literally “Which ball covers the largest distance”, after being explained in the introduction that the samples were of equally long duration but the balls were rolling at different speeds. Heller asked her participants “which ball was moving faster” which is almost the same as Houben who asked “which ball is rolling faster”. As can be seen in Figure 7.12 Heller’s data falls right on top of our psychometric function.

To resolve these questions, difference ratings could be gathered from participants using Houben’s sounds. An MDS study could resolve if people could indeed detect smaller differences in these sounds. Another possibility could be to find an acoustic description of the information that results in these detected differences. One possible technique could be kernel-based principal component analysis [77]. This is, however, a new method and its effectiveness remains to be tested.

7.4 Experiment 2: Absolute magnitude estimation

7.4.1 Method

The sounds of the second experiment were selected from the same set of recordings as used in the first experiment. The sounds from the 12 mm plate were used. The participants were not asked to judge the thickness in view of the difficulty the participants had with thickness discrimination, as shown in Figure 7.6. The experimental methodology for this experiment was absolute magnitude estimation, i.e., the participants had to indicate numerically their judgment of the diameter of the ball or the distance the ball traveled in the course of rolling. They were free to answer in any format they liked, that is using a number with a comma or dot if they wished, and using m, dm, cm or mm as units. Their answers were checked for validity, i.e., to be numerical and having a unit. If their response was not valid, they were asked to correct this, otherwise the next sound was presented. The experiment consisted of two parts, one for size estimations, the other for the speed estimations. Before the first part there were eight practice sounds for the participants to get accustomed to the range of the sounds and the method of

responding. Before the second part there were four such practice sounds. In between the two parts the participants were allowed to have a break if they wished. The experiments lasted about 20 min and the participants were paid 3.50 euros for participating. Twenty individuals participated. They were recruited at the university and were mainly students. All participants reported that they had no hearing problems.

In the first part, the participants were asked to judge the diameter of the ball producing the rolling sound presented to them. Two recordings of all balls with five different sizes and four different speeds were presented twice in random order, giving a total of four responses for each of the 20 combinations of speeds and sizes. In the second part, the participants were asked to judge the distance the ball covered in the course of rolling by listening to its recorded sound. Exactly the same sounds were used, but in a different random order.

7.4.2 Results

To check if participants were able to perform the task, we first tried to identify the relation between the physical size and the speed of the rolling ball in the recording, and the perceived size and speed. These physical and perceived sizes were log-transformed in accordance with Steven's law, which results in homoscedastic data, i.e., the variances of the logarithm of the perceived sizes were independent of the physical sizes, and the variances of the logarithm of the perceived speeds were independent of the physical speed. Using linear regression, a line was fitted through the average of the four answers of each participant.

$$\log(\text{perceived size}) = \text{factor}_{\text{size}} \times \log(\text{real size}) + \text{intercept}_{\text{size}}. \quad (7.1)$$

For the $\text{factor}_{\text{size}}$ we found values varying over the participants from 0.68 to 2.87 with an average of 1.71, and for the $\text{intercept}_{\text{size}}$ the values ranged from -13.0 to -6.08 with an average of -9.69.

These results are shown in Figure 7.13. The linear regression line fitted through all data is compared with the interquartile ranges of the data per size. The interquartile range is quite high, due to the fact that there is a large variation between participants. Individual participants often show much better performance, in the sense that they consistently give higher ratings to balls with a larger size. On average, the participants overestimated the differences in size resulting in a regression line with a negative intercept.

If we compare the predicted values of this linear regression with the averages of the four responses the participants gave for the four identical representations, we come to an explained variance of 44% when we apply the linear regression to the pooled data. When we apply the linear regression analysis to each participant separately, the explained variance increases to 79%. These individual regression lines all have positive values for $\text{factor}_{\text{size}}$. Hence, the participants were able to detect the differences in size.

Applying the same procedure to the speed of the ball, with

$$\log(\text{perceived speed}) = \text{factor}_{\text{speed}} \times \log(\text{real speed}) + \text{intercept}_{\text{speed}}, \quad (7.2)$$

we find for one participant a negative $\text{factor}_{\text{speed}}$ (-0.14). This means that this person judged faster balls as slower or, in other words, reversed labels, as did three of the 40 participants in the previous paired-comparison experiment in section 7.2. Because we cannot simply inverse the answers as in the previous paired-comparison experiment, we omitted this participant for the remainder of the analysis. For the other participants we found values for $\text{factor}_{\text{speed}}$ ranging from 0.023 to 0.83 with an average of 0.13. For $\text{intercept}_{\text{speed}}$ values were found from 0.025 to 0.23 with an average of 0.10. This time all participants underestimated the differences in the speed of the ball, and the intercept is positive. Besides the one reverse labeler, the result of one other participant was not included in this linear regression analysis, the participant with the highest $\text{factor}_{\text{speed}}$. Since the $\text{factor}_{\text{speed}}$ for this participant was 0.83 and the second highest $\text{factor}_{\text{speed}}$ was 0.29, with an average of 0.13, this high value is clearly an outlier in statistical sense. Remarkably, among all 20 participants this participant in fact best estimated the actual speed of the ball, since his or her $\text{factor}_{\text{speed}}$, 0.83, was closest to 1. For the other participants the $\text{factor}_{\text{speed}}$ was much lower.

So the data from those two participants was excluded from the analysis. The regression line through the data of the remaining 18 participants is shown in Figure 7.14. The slope of the line is close to the slope of the median of the data, including outliers. The fact that $\text{factor}_{\text{speed}}$ is much lower than unity means that, averaged over all participants, the listeners strongly underestimate the speed differences. If we compare the predicted values using this linear regression with the averages of the four responses the participants gave for the four identical representations we come to a, statistically not significant, explained variance of only 5% for the pooled data. If we, however, apply the linear regression per participant the explained variance is 39%, which is statistically significant. The low explained variance is no surprise after analyzing Figure 7.14, we see that the vari-

ation in perceived speed explained by the almost flat line is much lower than the variation indicated by the quartile ranges of the data. The much higher explained variance for the fit per participant, however, indicates that the participants are much more consistent in rating the faster ball as faster. This means that the participants were able to detect the differences in speed, but they underestimate the actual speeds and there were large individual differences.

Although this is the standard way to analyze data from an absolute magnitude estimation experiment, in our case this does not seem to be the optimal method. These plots in Figures 7.14 and 7.13 do not show that, although the variation between the participants was high, the participants themselves were generally reasonably consistent in their answers. Only by averaging across a large number of participants we get the shown response. Only few participants answered close to these average responses.

The main problem is that the answers of one participant cannot be compared with the answers of others directly, because each of them seems to have its own unit of measurement. By normalizing the data per participant we can only partially overcome this problem. In the stochastic estimation model the perceived magnitudes differ due to different perceived magnitudes but also due to a discriminational dispersion (see appendix A). We should only normalize for the scale differences in the perceived magnitudes, but we cannot distinguish this from the discriminational dispersion.

In the next section we will describe a better and more precise way to quantify the performance of the participants in this task. This method compares only responses from one participant at a time, and it makes thereby a comparison on one scale only. Via these comparisons within data from one participant, probabilities are calculated, and these probabilities can be compared across participants.

7.4.3 Psychometric functions

To calculate a psychometric function of the correct identification of the larger or faster of the two balls, we need data comparing balls of two sizes or of two speeds. We calculate these data in the following way. It is assumed that the participant estimates the size or speed of the two balls separately and then compares the two estimated values and responds accordingly. For each individual size or speed, we have four estimated values, which we treat as equally likely to be the result of the estimation process. Thus we have a total of sixteen equally likely comparisons, in which we can identify a correct or incorrect response. A third possibility is that the two estimates are equal. By counting the correct responses and giving two




FIGURE 7.13: Estimations for the size of the ball. Each circles indicates the medians across all data of 16 estimations by 20 participants; the bars indicate the interquartile ranges. The continuous line is the linear regression line through all data. The dashed line is the identity relation between physical parameters and the estimates.




FIGURE 7.14: Estimations for the distance traveled by the ball. Each circles indicates the medians across all data of 20 estimations by 20 subjects; the bars indicate the interquartile ranges. The continuous line is the regression line through all data. The dashed line is the identity relation between physical parameters and the estimates.

equal estimates half probability of being a correct answer and dividing by the total number of comparisons, in this case 16, we find what we will call a *derived percentage correct*. An example of how to calculate the derived percentage correct is given in Table 7.2. The participants were free to choose any step to round their responses to. The derived percentage correct represents the probability that the larger ball was indeed judged larger than the smaller ball. After calculating the derived percentage correct for all pairs, we can proceed to calculate the psychometric function as in the first experiment. However, the psychometric functions are now symmetric, because we have no order in which the two sounds were presented to the listener.

TABLE 7.2: Calculating paired data from absolute magnitude data, or the chance of a correct response in a paired comparison that we derive from two series of answers in an absolute magnitude estimation task. In the first row and column the physical sizes of the two balls are displayed. The second row and column display the four estimations of the listener. The table values code the comparison for each pair in the following way: 0: second estimate is larger; 1: first estimate is larger; .5: both estimates are equal. The average value of the table values gives the estimated proportion correct paired comparisons. The shown values are fictive and serve as an example, and would lead to a derived proportion correct of $12.5/16 = 0.78$ in which the first ball is judged larger.

		Actual ball 2 size [mm]:				
		25				
		Estimated ball 2 size [mm]:				
			28	31	39	43
Actual ball 1 size 45 [mm]:	Estimated ball 1 size [mm]:	35	1	1	0	0
		39	1	1	.5	0
		45	1	1	1	1
		52	1	1	1	1

The method has the advantage that it uses no reference and no measure of difference between the two sounds. The lack of reference is important because we know that a considerable part of our participants poorly recognize the absolute sizes of the balls. In Section 7.4.2 we saw that a linear regression on the pooled data results in a much lower explained variance than when a linear regression




FIGURE 7.15: The psychometric curve for the estimated size of the ball while the speed was constant. The data came from an absolute magnitude estimation experiment. The steepness β is indicated and the numbers between brackets present the 95% two-sided confidence intervals.

was applied for each participant separately. The lack of a measure of difference is important because the perceptual difference is the basis of the discriminability that we are trying to measure. If we would base our values on the mean or the median of the four responses we would have to convert the proportion of those two means into a probability, for which there is no known algorithm.

In the current case it is much less clear that the difference scale is logarithmic. In fact, due to the more scattered data, the fits are equally good when done on a linear scale. In order to compare the results from this paired-comparison experiment with the previous absolute magnitude estimation experiment, we will use a logarithmic scale here, too.

In Figure 7.15 and 7.16 we see the results. Again, the psychometric function for size, Figure 7.15, is much steeper than the one for speed, Figure 7.16. The fit is quite good resulting in small confidence intervals. Because of this good fit we can also try to use only a subset of the data to fit the functions. In Figure 7.18 the steepness of the psychometric function, β for judging travelled distance, is plotted as a function of the diameter of the ball. The different ball sizes are not averaged as was the case in Figure 7.16, but analyzed separately. Similarly, in Figure 7.17 the size data were used and analyzed for each ball speed separately. In these figures, we can see that there is some statistically significant variation in the steepness of the psychometric function, β . Speed judgments seem easiest for the smallest ball, which had a diameter of 25 mm. Remarkably, the next smallest




FIGURE 7.16: The psychometric function for the estimated traveled distance of the ball, while the size was constant. The data came from an absolute magnitude estimation experiment. The steepness β is indicated and the numbers between brackets present the 95% two-sided confidence intervals.




FIGURE 7.17: The steepness of the psychometric function for the diameter judgment for separate speeds. The data came from an absolute magnitude estimation experiment. The error bars denote the two-sided 95% confidence intervals.




FIGURE 7.18: The steepness of the psychometric function for the traveled distance judgment for separate sizes. The data came from an absolute magnitude estimation experiment. The error bars denote the two-sided 95% confidence intervals.

ball, that of 35 mm diameter, is the most difficult to judge. If there were a simple interaction of the ball size on the judgment of the speed of the ball, we would expect a monotonic change of the steepness for speed judgment when the size of the ball increases. Similarly these figures show no systematic interaction of ball speed on judgment of the size. In the current data, however, we can see that the capability to discriminate speed does not increase or decrease for larger sizes, it is merely so that for some of the balls, specifically the 25 mm one, it was easier to detect the speed than for other ones, specifically the 35 mm one.

7.5 Discussion

In both experiments, the participants were able to detect differences in the sizes and in the speeds of a rolling ball. Participants were sometimes able, but generally performed very poorly in estimating the thickness of the plate on which the ball rolled in experiment 1. This is somewhat surprising because in other experiments, discussed in Chapter 8, we established that the participants could hear the differences between the sound of a ball rolling over a thick plate and the sound of a ball rolling over a thin plate. Also, in daily activities people often tap on a piece of wood of, for instance, a table to hear if the top plate is thin, thick or hollow. This surprisingly does not extend to rolling sounds and more specifically, they fail to use the differences they hear to estimate the thickness of the plate.

An absolute size experiment was also done by Grassi [31]. He asked participants if they could estimate the size of a bouncing ball. Grassi concluded that listeners underestimate the size of the ball. From the figures plotted in his paper we can also conclude that his participants overestimate the size differences, but have a negative intercept. In his experiment six out of seven balls were underestimated in size. This corresponds very well with our findings in the size estimation using absolute magnitude estimation.

By plotting the percentage correct as a function of the difference of the logarithm of the ball sizes, data were obtained through which a well-defined psychometric function could be fitted. From the fact that the fit is very good in the case that we have the best data (see Figure 7.15) we can derive a distance measure that we can also use when the data are somewhat more scattered (e.g., see Figure 7.9). The psychometric functions shown in this chapter are based on data from all participants. In this way we can calculate the estimated percentage correct, EP1, that a listener will hear the difference in size or speed between two of the balls:

$$\text{EP1} = \text{gaussian}(\log(s_1) - \log(s_2), 1/\beta), \quad (7.3)$$

where the two sizes or the two speeds are indicated with s_1 and s_2 . The cumulative Gaussian distribution is used:

$$\text{gaussian}(x, \sigma) = \frac{1}{2} \left(1 + \text{erf} \frac{x}{\sigma\sqrt{2}} \right). \quad (7.4)$$

We can calculate the size or speed difference, or more precisely the ratio, s_1/s_2 necessary to realize that more than EP1 of the participants judge the correct ball as being larger or faster:

$$\frac{s_1}{s_2} = \exp \left(\frac{\sqrt{2}}{\beta} \text{erf}^{-1}(2\text{EP1} - 1) \right), \quad (7.5)$$

where erf and erf^{-1} stand for the ordinary and inverse error function (see, e.g., Arfken and Weber [4]).

7.5.1 Absolute magnitude estimation versus paired comparison

Apart from these issues, the merits of both methods to generate the psychometric functions can be compared. In our design we had four speeds and five sizes leading to 20 different sounds. In the absolute magnitude estimation experiment we presented all sounds four times, with two different questions, one for the size

and one for the speed, which leads to 160 presentations to the 20 participants. We presented each pair once in each order to the 23 participants in a paired comparison task, which leads to 140 presented pairs. Thus, the total numbers of sound presentations are similar. Yet, the confidence intervals for the steepness of the psychometric functions resulting from the absolute magnitude scaling data are smaller. This is due to the advantage of absolute magnitude estimation, which is that there is more information gathered for each presented sound, namely a numeric value that we can compare with all the other responses by the same participant. In paired comparison only a small amount of information is gathered after presenting two sounds, either the first or second is estimated larger or faster. The psychometric functions can, therefore, be fitted with fewer hours of participant testing. The disadvantages of absolute magnitude estimation is that this task is harder for the participants, and we cannot tell immediately if an answer is right or wrong. Participants almost never give the exact answer, but this does not mean that the data points do not fit well on a psychometric curve. Another advantage of paired comparison is that we can choose to fit the psychometric curves symmetrically or asymmetrically. A disadvantage for paired comparison is that all sounds considered need to be presented in pairs. Especially when examining interaction effects, this will lead to excessively many pairs. We dealt with this problem by only presenting a small subset of the pairs of sounds in which both speed and size differed simultaneously, presenting all these pairs twice, once in each order.

7.5.2 Interaction

We have looked for interaction effects between size and speed judgments in various ways. Using the data obtained in the absolute magnitude estimation, a linear regression model was used to fit the data of the participants. Extending the linear model of Equations (7.1) and (7.2), the mean of the four estimations was fitted for each participant individually with the following function:

$$\begin{aligned} \log(\text{perceived size}) = & \text{factor}_{\text{size}} \times \log(\text{real size}) + \text{intercept}_{\text{size}} + \\ & \text{factor}_{\text{covary}} \times \log(\text{real speed}) + \\ & \text{factor}_{\text{inter}} \times \log(\text{real size}) \times \log(\text{real speed}). \end{aligned} \quad (7.6)$$

The addition of one or both of the two terms, $\text{factor}_{\text{covary}} \times \log(\text{real speed})$ or $\text{factor}_{\text{inter}} \times \log(\text{real size}) \times \log(\text{real speed})$, did not increase significantly the explained variance when such a model is fitted through the average estimates of all

TABLE 7.3: The steepness β of the psychometric curve can be calculated from part of the data where the size as well as the speed differed. The performance in the absolute magnitude task is split into three categories: In the first column the results for half of the data, where the largest ball is the fastest, are used to calculate the steepness, β . In the third column the other half of the data is used to calculate the steepness, thus where the smallest ball is the fastest. In the second column all data is used, and this consists of an equal number of cases where the smallest and the largest ball is the fastest.

	largest is fastest	counter balanced	smallest is fastest
distance	(0.56-) 0.66 (-0.76)	(0.76-) 0.82 (-0.88)	(0.95-) 1.0 (-1.1)
diameter	(1.9-) 2.0 (-2.1)	(2.0-) 2.2 (-2.3)	(2.2-) 2.5 (-2.7)

data of all participants. For the speed estimation the explained variance increased from 5.3% to 6.0%, and for the size estimation it remained at 44%. When such a plane is fitted for each participant individually, however, the explained variance increases with the inclusion of these two parameters. For the speed estimation the explained variance increased from 39% to 72% and for the size estimation it increased from 79% to 88%. This leads to the conclusion that there is no systematic interaction effect present in these data. Apparently the interaction term from one participant is canceled by that from another. We can, however, only compare the regression coefficients of the fits and not the accuracy of the individual estimations of speeds or sizes because of the large individual differences between participants.

After constructing the psychometric functions, we showed in Figure 7.17 that the steepness of the psychometric function for the speed judgment varies for different ball sizes and, in Figure 7.18, that the steepness of the psychometric function for the size judgment varies for different ball speeds. There seems, however, no general trend in the sense that the speed of larger balls is easier or harder to discriminate than the speed of smaller balls. It seems merely that the speed of some of the balls is easier to discriminate than the speed of other balls. We can only guess for the cause of this difference, but there are large temporal differences in the sounds of the rolling balls. Sometimes, especially for smaller and faster balls, the balls bounce more readily. In the absence of bouncing, the sound of certain balls at certain speeds shows a more pronounced temporal variation in amplitude. This amplitude modulation corresponds with the rotational velocity of the ball, and has a strong influence on speed and size estimations [40]. If listeners are capable of interpreting these extra sources of information, and they are for amplitude modulation [40], their availability could explain this part of our data.

TABLE 7.4: A summary of the steepness parameter β as found in the different parts of this paper. The two methods are magnitude estimation and paired comparison

	magnitude estimation	paired comparison
diameter:		
equal distance	(2.2-) 2.4 (-2.5)	(2.1-) 2.6 (-3.1)
distance:		
diameter equal	(0.80-) 0.86 (-0.92)	(1.5-) 1.7 (-1.9)
diameter different	(0.76-) 0.82 (-0.88)	(0.89-) 1.2 (-1.4)

Furthermore, there could be a subtle difference between the ways participants listen to sounds when presented one at a time or when presented in pairs. The fact that people listen in different ways to the same kind of sounds is not new. A classic way is Gaver's [24] division in musical and everyday listening. For our case we could hypothesize two models in which the listener responds to a sound pair. In the first model, the listener would treat each sound separately, and come to an internal representation of the size or speed of the ball. In paired comparison, these two internal representations are compared to come up with an answer. This will lead to exactly the same results compared to an absolute magnitude estimation task in which the participants will answer directly according to their internal representations of the speeds or sizes of the balls. In contrast, when the participants compare the two sounds on a somewhat lower processing level, e.g., by listening to temporal patterns or the brightnesses, then there could be a difference in both methods. Assume that people have some short-term memory space available to remember, for instance, the temporal pattern over a short time, and that the capability to remember this pattern fades over time. In this case, for a paired comparison task the two sounds are compared at a higher resolution than in an absolute magnitude estimation task. This will lead to steeper psychometric curves.

Considering the possibility of comparing sounds at different levels, it is interesting to see that the psychometric functions for the size estimations are comparable for the two methods whereas the functions for speed estimations differ. On the basis of the foregoing discussion, this could mean that listeners are able to judge the size differences of the two balls after they have built the complete internal representation of the size of the rolling balls. With this we mean that the listener has an idea of the actual diameter for both balls, and compares these two quantitatively. This explains why, in the size judgment task, the steepness of the results obtained with absolute magnitude estimation is equal to that of the paired

comparison task, see Table 7.4 for a summary of the β values derived in the different experiments of this chapter. In the speed judgment task, however, listeners perform better in the paired-comparison task than in the absolute magnitude estimation. Again, following the discussion of the previous paragraph, this would imply that they do not build the full internal representation of the speeds of the balls, but, in a paired-comparison task, they compare the sounds in a more direct manner. This more direct method could consist of comparing lower-level perceptual attributes of the auditory signals such as the brightness or roughness of the sounds and lead to a higher sensitivity than in the magnitude estimation task. For a definition of brightness and roughness see [112].

In the cases where both the speeds and the sizes of the balls differ, listeners can no longer rely on the lower level properties of the auditory signals, because they vary with both the speed and the size of the ball. They, therefore, compare both sounds at the level of the internal representation of the ball size, and the steepness of this psychometric function for this paired comparison is more equal to the one constructed from the absolute magnitude data, and much flatter than the psychometric function for this paired comparison where the speed of the ball is not varied. For the speed estimation the situation is slightly different. The listeners can still make comparisons on the basis of the lower-level perceptual attributes of the auditory signals, and compensate for the ball diameter, for which they can make a much better estimate. This explains the higher steepness for the distance estimates, using the paired comparison case in which the balls are of different diameters, compared to the absolute magnitude tests.

7.A Psychological scales

In this appendix, we will analyze the perception of the size and speed of rolling balls using a concept that was introduced by Thurstone called “discriminal dispersion” as introduced in [102]. Our reasoning will closely follow that of Thurstone, see also [103] and [25], but confines it to our particular case and to the psychometric functions. The essence of this concept is that the perception of a stimulus can be represented by a random variable. The actual or physical stimulus value, in our case the size or speed of a rolling ball, is represented by a random variable with mean value $\bar{S}$, on the relevant psychological scale. For each individual representation of a stimulus there is, for various reasons, a random deviation from this mean value.

We will assume that the variance of this random variable remains constant for all ball sizes or speeds, although this is no necessity for Thurstone’s theory, but it simplifies our calculations. Furthermore, we will assume, as is essential in Thurstone’s theory, that listeners perceive the size or speed differences of the two balls by subtracting the internal representations of the size or speed of the balls indicated by the two random variables S_1 and S_2 . The result is a new random variable S_d with $E[S_d] = E[S_1] - E[S_2]$ and $\text{var}[S_d] = \text{var}[S_1] + \text{var}[S_2]$. When this new random variable is positive, this means that the first presentation is judged larger or faster. This means that the shaded proportion, shown in Figure 7.19, represents the probability $P(S_{d1,2} > 0)$ that the first presentation is judged larger or faster. This probability $P(S_d > 0)$ is a function of the difference $\bar{S}_2 - \bar{S}_1$. This means that for another pair of quantities on the same scale, say S_3 and S_4 , that have a different position on the same psychological scale, but have the same distance, this leads to the same $P(S_{d1,2} > 0) = P(S_{d3,4} > 0)$; it also means that $P(S_{d1,3} > 0)$ can be calculated when knowing $P(S_{d1,2} > 0)$ and $P(S_{d2,3} > 0)$. Such calculations are only permitted if these probabilities form a metric space.

If the probability $P(S_{dx,y} > 0)$ with x, y indicating a numeric index for the two sounds, forms a metric space, there should be a distance $d(x, y)$ that conforms the prerequisites for a metric space: non-negativity of $\forall_{x,y} d(x, y) \geq 0$, identity $d(x, y) = 0$ iff $x = y$, symmetry $\forall_{x,y} d(x, y) = d(y, x)$, and the triangle inequality $d(x, z) \leq d(x, y) + d(y, z)$. A valid distance for our space is:

$$d(x, y) = |P(S_{dx,y} > 0) - 1/2|$$

All together, in such a metric space we can calculate the scale values from the frequency table, and this is exactly what is done by Thurstone’s method.




FIGURE 7.19: (a) Distributions of two perceived quantities on a psychological scale. (b) The distribution of the discriminational differences on a psychological scale. The shaded portion to the right of the zero point the probability that the investigated magnitude of stimulus 1 is perceived larger than magnitude of stimulus 2.

Using Thurstone's method on our frequency tables, Table 7.1 and 7.2 of the preceding chapter, results in a number of relative scale values for the perceived size and speed of the ball. In Figures 7.20 and 7.21 we plotted these scale values plotted against the logarithm of the physical or actual parameters. It can be seen that the data from the absolute magnitude data experiment follow the logarithmic scale more closely than the data from the paired-comparison experiment. Calling the physical value A and the position of the mean of the random variable on the psychological scale, $\bar{S}$, results in

$$\bar{S} = \alpha \ln(A), \quad (\text{A-1})$$

where α is some unknown constant. The probability $P(S_d > 0)$ is a function of the difference $\bar{S}_d = \bar{S}_2 - \bar{S}_1$ and, hence, also of $\alpha(\ln(A_2) - \ln(A_1))$. It represents the proportion of times S_2 is reported as coming from the larger or faster ball than S_1 . Thus, if the representation of this difference S_d is positive, the listener perceives the first sound as larger or faster. The probability $P(S_d > 0 | S_2 - S_1)$ or $P(S_d > 0 | \alpha(\ln(A_2) - \ln(A_1)))$ results in the psychometric function. If we assume that the distribution of the random variable is Gaussian, the psychometric functions will have the form of a cumulative Gaussian, and this is what we chose for the psychometric functions in this chapter.

Following Stevens' law, Equation (A-1) should contain an extra constant, indicating the proportionality of the relation, $\bar{S} = \alpha \ln(A) + \ln(c)$, which was included in the linear regression in section 7.4.2. Since we are now only interested in the difference of two scale values, these constant proportionality factors would cancel each other and, therefore, we do not consider them in this appendix. Hence, we can use a single parameter to describe the psychometric function, $\beta = 1/\alpha$,




FIGURE 7.20: Relative perceptual scale values for different ball sizes, derived using Thurstone's method. Data are from a paired comparison experiment (x-marks) and an absolute magnitude scaling experiment (squares). Positions are relative to the smallest ball. The line is fitted from the data under the assumption that Stevens' law holds, i.e. two stimuli are equally different if their ratio is the same, hence the logarithmic scale for the horizontal axis.




FIGURE 7.21: Relative perceptual scale values for different ball speeds, derived using Thurstone's method. Data are from a paired comparison experiment (x-marks) and an absolute magnitude scaling experiment (squares). Positions are relative to the ball with the lowest speed. The line is fitted from the data under the assumption that Stevens' law holds, i.e. two stimuli are equally different if their ratio is the same, hence the logarithmic scale for the horizontal axis.

and we can calculate the relative position of the ball size on the psychological scale in units of σ , the standard deviation of the random variable that represents the perceived size or speed of the ball. This means that if the distances on the (logarithmic) psychometric scale are twice as large, this leads to β being twice as large and thus to twice as steep a psychometric function.

The method of Thurstone has been criticized, see for instance Greenwood and Williamson [32]. The main points of critique come from the fact that it is an “indirect” measurement technique, i.e., the scale is not measured directly but derived from measurement results. For instance, one criticism is that it measures the scale in some region and not a global scale. But the same can be said of magnitude estimation, which is supposed to be direct. In the end we are searching for the mental representation of the objects heard, and for that no direct measurement techniques exist.

The difference between Thurstone’s method and the method based on the psychometric functions that are estimated in the main part of this chapter, is that Thurstone’s method estimates the scale values from a frequency table, and the psychometric functions estimate the accuracy of the perceived scale from the frequency table and known scale values. We know what the sizes and the speeds of the balls were when we recorded the sounds. Using Thurstone’s method, we can show that Stevens’ law holds for the data generated by the absolute magnitude experiment, by plotting the relative distances of the stimuli on the perceptual scale, as found by using Thurstone’s method versus the logarithm of their sizes or speeds. This is shown in figures 7.20 and 7.21. From the good fit of the straight line, we conclude that Stevens’ law indeed holds, and we can safely use a log-ratio scale on the psychometric functions.

8 | Using Multidimensional Scaling to Determine Distances in the Perceptual Space of Rolling Sounds¹

Abstract

The experiments described in this chapter were performed to allow the construction of a perceptual space for rolling sounds. The same recorded sounds as in Chapter 7 of this thesis were used as experimental stimuli. All sounds were combined in pairs and participants were asked to quantify the perceptual difference between the two sounds of a pair on a five-point scale. Using this method, all possible perceptual differences could contribute to the difference scalings without a need for the participants to interpret these differences in terms of underlying object properties such as size or speed. Using a multidimensional scaling algorithm, the scaled differences were used to construct a perceptual space of these sounds. A good fit was obtained for a space having three dimensions. None of these dimensions corresponded directly with one of the three mechanical parameters varied in the rolling sounds. Differences in speed led, for the parameter ranges used, to a smaller distance in the constructed perceptual space than differences in size. This corresponds to results from direct scaling experiments using the same sounds. Remarkably, differences in the plate thickness resulted in large distances in the perceptual space, while this mechanical property was poorly recognized when participants were directly asked to estimate it. In the perceptual space, the direction that correlated best with plate thickness was almost perpendicular to the directions that correlate best with the speed and size of the ball. This orthogonality might indicate a perceptual process that separates the more relevant auditory information from the less relevant.

¹Part of this work was presented at the ASA Providence meeting in 2006 [94]

8.1 Introduction

Up to now, all our research on the perception of rolling sounds has been focused on determining the auditory capability of humans to detect some mechanical feature of the rolling sound. In this study, we tried to quantify the differences listeners can hear in a rolling sound without labeling what the differences are. Thus, after listeners were presented with two sounds they were asked to respond how different they thought the sounds were, whereas in previous research they had to answer some specific question, such as which of the two sounds represented the bigger or faster rolling ball. There are two important differences between those two methodologies. When people can answer some specific question, like “which of the two balls is the larger one” correctly, this requires two skills of the listener. The listener has to be able to detect the differences in the sound caused by this change in size of the ball, and second, he or she has to be able to interpret these differences. The listener has to separate the differences in the sound caused by this change in size from other changes in the sound, and this will enable the listener to derive the relative sizes of the balls. Only if both detection and interpretation have taken place, people will come to the correct answer to such questions. This was tested in previous research. In the current study there is no specific question, and people do not have to interpret the differences they hear, and we can concentrate on the detection of the differences in the sounds.

Next, we would like to figure out whether people can “tune their ears” to listen to the size of rolling balls and then use a different tuning when listening to the speed of the ball. In other words: Does the question posed to the listener have an influence on the perception process. We will try to answer this question by looking at a configuration where this adaptation or “tuning of the ears” does certainly not take place, because the question asked to the listener is always the same. If we then do not find enough differences to explain the differences found in Chapter 7, we may conclude that this tuning of the perception process indeed takes place. This tuning is similar to schema-driven perception where different schemas can be activated. For an explanation on schema-driven perception see [7][106][67]. Note that the participants were, in the papers mentioned before, asked directly to judge the size or the speed of the ball, and in one case the thickness of the plate of the recorded rolling sounds. For the current work we asked the participants to judge the difference they perceived between the sounds without identifying the origin of these differences. A multidimensional scaling was then applied to determine the underlying perceptual space of these rolling sounds.

Multidimensional scaling has been applied earlier in sound-perception research. For instance, it was used to study differences for signal processing techniques that influence the overall differences in sounds [3]. The stimuli in that study were synthetic; pure tones with or without applied frequency modulations, noise bands with or without amplitude modulation and harmonic complexes, whereas we used recorded sounds. Also, these authors focused on the differences in scaling between adults and 7 or 10 year old school kids. By correlating the periodicity and temporal fine structure of the sounds with the found MDS axes, they were able to determine that adults and many children appear to base their classification on both the temporal and the spectral structure of the sound.

In an attempt to evaluate the perceptual relationships between sounds from 16 musical instruments, Grey [33] used stimuli, the synthesis of which was based upon a spectro-temporal analysis of actual instrument tones. Using this analysis, he tried to explain the confusions listeners made between instruments. This idea was evaluated further by McAdams [67, Ch. 6] who used multidimensional scaling in an attempt to explain the recognition of sound sources and events.

Some approaches using multidimensional scaling, quite similar to the work presented here, were aimed at studying the processes that enable listeners to detect mechanical properties of sound generating systems. The psychomechanics of simulated impacted mallets was studied by McAdams, Chaigne and Rousarie [68]. Based on multidimensional scaling, a two-dimensional perceptual space was constructed, and signal parameters were found that explained a significant amount of variance in the response data. The two best descriptors appeared to be the long-term decay constant and the spectral center of gravity. In another work by Canévet, Habault, Meunier and Demirjian [10], the auditory perception of sounds radiated by a plate excited by a transient point source was studied. In this work too, synthesized sounds were used. In their multidimensional scaling they found three dimensions. They found that a first dimension was related to the “structural damping and high-frequency components”. Another dimension was related to the “tonal character of the resonant part of the sound”. The third dimension was related to the “sharpness of the transient part of the sound”.

Note that in these studies fast decaying sounds were used that were synthesized and relatively short. Correlations with acoustic signal parameters were studied. For the current work we used recorded sounds that are relatively long and steady-state, and we will try to correlate the results from the multidimensional scaling with the mechanical parameters of the sound-generating system.

Then we will discuss the auditory processing mechanism that enables listeners to detect the mechanical properties of the sound-generating system.

8.2 Stimuli

The stimuli consisted of recordings of wooden balls rolling over wooden plates. In the recording of these sounds, the ball diameter, the rolling speed and the thickness of the plate were varied. The balls and plates were chosen to give good rolling sounds and to be the same or similar as used in earlier experiments [40]. The balls were made of beech wood and had diameters of 25, 45 and 68 mm. Due to the production process, the balls were not perfectly spherical, leading to temporal variations in the sound, such as amplitude modulation, or even some bouncing for the smallest and fastest balls.

In order to give the balls a well defined velocity along the center line of the plate, they were rolled from a slide onto the plate. The height of the slide was 25 cm. The balls were released from various heights on the slide to vary their velocity. In order to measure the velocity of the ball, its course interrupted the beams of six independent light-gates, placed at intervals of 20 cm. The measurements showed that, during one run, the ball's velocity was constant within a few percent. The velocities we used in the experiment were 0.48, 0.91, and 1.63 m/s.

The slide was separated from the plate by a narrow slit, about 0.5 mm in width, to avoid the transmission of vibrations from slide to plate. The slide had a smooth bend near its end to transform the vertical velocity of the balls into a horizontal one. At the other side of the plate, the ball was left free to roll off.

The plates, 50 cm wide and 120 cm long, consisted of MDF, a kind of wood-particles pressed and glued together. The thickness was 8, 12 or 18 mm. The advantage of this kind of material over real wood is that the material properties do not change with position and direction on the plate.

Seen from above, the microphone was placed 70 cm to the right of the long side of the plate, and halfway its length. It was placed 70 cm above the plate, pointing at it with an angle of 45 degrees. The closest distance between microphone and plate was 100 cm. The microphone itself was a small diaphragm cardioid condenser microphone, type Røde NT5. It was connected to a sound card, type Terratec MIC8, that recorded the sounds with 24-bit resolution at a sample frequency of 48 kHz on a standard PC.

From the middle of the sound, 600 ms intervals were cut out. The waveforms were gated on and off with half of an 8 ms Hanning window. The two sounds of

a pair were presented with a 250 ms silent interval. The sounds were presented to the participants via headphones, at their original relative sound levels.

Because spectral differences are often considered as a very important source of information, the spectral differences between the sounds of these mechanical systems were analyzed and the results can be found in Chapter 7.

8.3 Method

From the recorded sounds we selected all three ball sizes, three ball speeds and three plate thicknesses, resulting in 27 different sounds. All possible pairs were constructed except those where the two sounds were identical, resulting in $26 \times 27 = 702$ sound pairs. These sound pairs were presented in random order via headphones to the participants. Their task was to scale the differences between the sounds on a five-point scale ranging from “no audible differences”, scale value 1, to “very different”, scale value 5. The intermediate steps were not given a label. The main part of the experiment was preceded by eight practice sound pairs to let the participants get familiar with the interface. The main experimental part itself was split into four parts with a possibility of a break in between. The experiment lasted about 40 minutes and the 27 participants were awarded 5 euro for participating. They were recruited at the university and were mainly students. All participants reported to have no hearing problems.

8.4 Results

On the dissimilarity scale from 1 to 5, the average response across all participants and all sound pairs was 3.36 while the standard deviation was 1.01. The sound pair with the highest rated difference had an averaged response of 4.5, the lowest response was 1.55. In Figure 8.1 the confidence intervals for the means of the difference scalings given by the participants are presented. The two ratings given by each participant for each pair in the two orders were averaged. On this scale, the width of the confidence intervals ranges from 0.30 to 0.74.

In the following, we analyze the mechanical properties of sound pairs with the largest and the smallest differences. The ten sound pairs that were most alike according to the participants were rated from 1.6 to 1.9. The two sounds in each of these pairs differed in the speed of the ball, and four of them also differed in size. For those four pairs, that differed in two dimensions, the ball with the highest speed was the smallest in diameter. The sound pairs with the five highest




FIGURE 8.1: The confidence intervals of the difference ratings. For all sound pairs the 95% confidence intervals are plotted against the mean of the difference ratings. Due to the limited number of data, some values for the mean of the difference rating occur more than once, in which case the mean values are spread a little bit to make the individual lines visible.

rated differences differed in all three dimensions. Their ratings ranged from 4.4 to 4.5 on the same scale. Of the next five, all with rating 4.4, only one differed in three dimensions and two differed in plate thickness and the size of the ball, the other two differed in plate thickness and the speed of the ball. As we will see later, this simple analysis gives a good indication of the results obtained with the more sophisticated MDS analysis.

8.4.1 Multidimensional scaling

So far we only discussed the average dissimilarities of all 351 sound pairs. These dissimilarities represent distances between the points in the perceptual space that represent the different sounds. We can now try to find coordinates in a space in such a way that the distances between the points are proportional to the average dissimilarity as judged by the participants of our experiment. Depending on the dimensions of the space that we try to construct, there always remains a discrepancy between the distances in between the points in the constructed space and the average dissimilarity of the sounds as reported in the experiment. The mean squared average of this discrepancy over all sound pairs is called the stress. The stress depends on the number of dimensions of the underlying space, where the




FIGURE 8.2: The three panels represent the perceptual space of the sounds. In all three panels one of the three dimensions of the multidimensional scaling analysis is perpendicular to the paper, so we can see only the two remaining dimensions. The mechanical properties of the plate are represented in the symbols. The size of the symbols represents the thickness of the plate, where the largest symbol corresponds to the thickest plate. The type of the symbol represents the speed of the ball. The circle was the slowest, the square the middle and the diamond the fastest. The fill color of the symbol represents the size of the rolling ball, with the open symbol representing the small ball, the gray one the middle and the black one the largest balls.

representation of the stress as a function of the dimensions of this space is called the stress function. The stress function is always decreasing with increasing number of dimensions of the underlying space but typically shows a steeper decrease up to a certain number of dimensions. This number of dimensions is then regarded as the best choice to describe the perceptual space. In our case, however, the stress function was not conclusive about the required number of dimensions. We chose to use three dimensions because the results were easiest to interpret in this case. In addition, as we will see later, this enabled us to transform the axes resulting from the MDS analysis into new axes that correlate best with the three dimensions of the mechanical variations applied in our setup.

The software program used for the multidimensional scaling was XGms as described in [65]. The paradigm used is referred to as double-stimulus dissimilarity scaling. The distance measure used had a Minkowski norm of rank two, which means that the ordinary sum of squares of the residue was minimized.

Within this MDS space we can use linear regression techniques, to find the directions that correlate best with the variations in the mechanical parameters. The angles between the thus found directions, with maximum correlation with the thickness, speed and size are: 85 degrees between thickness and speed, 99 degrees between thickness and size and 130 degrees between speed and size. From this we can conclude that the thickness dimension is nearly perpendicular to the other two dimensions and that the speed and size dimensions interact, with a larger speed corresponding with a smaller size. It appears that a larger size compensates, perceptually, for a faster ball. Thus, two rolling sounds are perceived more different, when the largest ball is the slowest than when the largest ball is rolling fastest.

The three directions found in this way can also be used to form a basis for a transformed space to describe the perceptual positions of the sounds. This is, however, only possible if the three directions are independent of each other, that is, have components in along all three axes of the (original) MDS space. The directions in perceptual space which correlate best with the variations in the mechanical parameters do have indeed components along all three axes in the original MDS space and thus form a good basis for a transformed perceptual space. The position of each individual sound in terms of the new basis can be found by multiplying it with the transformational matrix. It should be emphasized that this transformation is not a simple rotation of the original three-dimensional MDS space, but that the transformational matrix is non-orthogonal. In the appendix to

TABLE 8.1: Correlations between the mechanical parameters of the objects and the original axes found by the MDS analysis.

mech. param.	mds axes		
	dim1	dim2	dim3
thickness	-0.21	0.84	0.39
speed	-0.56	-0.15	-0.02
size	0.64	-0.32	0.73

this chapter this transformation of the originally-derived MDS space is explained in more detail.

The explained variance of the parameters can be calculated by squaring the correlation coefficient between the parameter and the axis. The correlation coefficients are found in Table 8.1 for the original basis and in Table 8.2 for the transformed basis.

When the distance between two sounds in the MDS space is large, this means that the participants have rated the differences between the two sounds as large. When looking at the correlation coefficients in Table 8.2, the correlation coefficient for size, 0.80, is, in absolute sense, larger than the one for speed, -0.57, indicating that differences in the speeds of our sounds were given less weight than differences in the sizes of the balls. If the judgment of the participants is represented by the difference between the projection of these speeds on this axis, this corresponds with a much poorer performance than in the case of the size judgment. This is exactly what we found in the experiments in Chapter 7 where the participants were directly asked to judge the size and speed of a rolling ball.

However, the situation is remarkably different as far as plate thickness is concerned. Here, the correlation coefficient between mechanical parameter and the corresponding axis in perceptual space is, in absolute sense, the highest, 0.94. We can, therefore conclude that the listeners are very well able to hear differences in the rolling sounds associated with the plate thickness. Since in Chapter 7 it was shown, however, that listeners perform badly in thickness judgments, we can conclude that the listeners do not use the perceptually available information in the thickness judgments.




FIGURE 8.3: The three plots represent the perceptual space of the sounds after applying the basis transformation. The axes of the constructed three-dimensional space now indicate the mechanical parameters that have the highest correlations with these axes. In all three panels one of the three dimensions of the MDS analysis is perpendicular to the paper, so we can see only the two remaining dimensions. The mechanical properties of the plate are represented in the symbols. The size of the symbols represents the thickness of the plate, where the largest symbol corresponds to the thickest plate. The type of the symbol represents the speed of the ball. The circle was the slowest, the square the middle and the diamond the fastest. The fill color of the symbol represents the size of the rolling ball, with the open symbol representing the small ball, the grey one the middle and the black ones the largest balls.

TABLE 8.2: Correlations between the mechanical parameters of the objects and the transformed axes, as reconstructed with the linear regression method.

mech. param.	mds axes		
	dim1 (thickness)	dim2 (speed)	dim3 (size)
thickness	0.94	0.00	0.00
speed	0.00	0.56	0.00
size	0.00	0.00	0.82

8.5 Discussion

The thickness of the plate has a pronounced effect on the sound generated by a ball rolling over a plate. Due to the cut-off effect the sound generated by the plate is high-pass filtered, as the lowest frequencies simply do not radiate from the plate. This so-called cut-off frequency is lower for the thicker plates than it is for the thinner plates. On the other hand, the modal frequencies are higher for the thicker plates than they are for the lower plates. The third effect is that the bouncing process of the ball is altered because the thicker plate is stiffer than the thinner plate. Finally, the frequency dependent damping parameters vary with the plate thickness [12]. These variations do not induce simple alterations in the spectrum but bring about changes that induce differences in almost any perceptual attribute of the sound, such as loudness, brightness and roughness.

Despite all these complex alterations in the spectrum, the reconstructed basis shows that the ball size and speed variations result in audible differences that are perpendicular to the audible differences caused by the plate thickness. This is highly remarkable because we tend to believe that spectral content of the sounds plays a very important role in size and speed perception. Another very remarkable result of the analysis is that the thickness of the plate is best represented in the perceptual space whereas it is not recognized when participants are directly asked which of two sounds represents a ball rolling over the thicker plate, as was discussed in Chapter 7.

The fact that the plate thickness is represented as a separated dimension almost perpendicular to the others seems illogic at first sight. Perhaps compatible results may have been found by Freed [22] who noted that listeners were able to separate out interacting dimensions of a stimulus in order to make an

environment-oriented judgment. In his work the perceived mallet hardness was independent of the pan struck with it.

The auditory information is, so to say, pre-processed to exclude the plate thickness information from the system. But, had the plate thickness not been nearly orthogonal to the two other dimensions, then it would not have been possible to ignore this information. It appears that the size and speed of the ball are perceptually more relevant and that these two are separated from other causes of variations in the sound, thus enabling perceiving them separately. The fact that the thickness is nearly perpendicular to the size and speed in the generated difference scale means that if we simply ignore this axis in the perceptual space, the perception of the size and the speed is not influenced by the differences caused by the plate thickness. We may perhaps also compare this with the process that separates room acoustics from sound generating events; on every location in the room a constant filter is imposed over any sound that we hear, due to the response of the room. This filter can be absent outdoors, but can also cause very severe alteration of the spectrum of the sound when, e.g., considering small rooms with reflecting walls. Yet, when we try to recognize someone's voice, we can do this both outdoors and in a small reflecting room. We are apparently able to ignore the room acoustics and focus on the sound source. Thus, we are able to separate the properties of the sound source, the voice, from that of the path of the sound to our ears, i.e., the room which acts as a filter on the sound. It could well be that the same process underlies the inability to estimate the thickness of the plate. It may be considered part of the path from the sound source to our ears. The effect of the plate is then similar to that of the room. It is a, nearly constant, filter with some crests and peaks that are determined by the thickness of the plate. The only difference is that the plate is part of the sound generating system, so that there cannot be sound without the plate. Thus we cannot listen to the rolling sound without the filter, whereas we can listen to somebody's voice without the filter of the room, namely outside.

Another effect that may be explained by this reasoning is the inability of participants to detect whether a ball is rolling from or towards the edge of a plate, which we tested and reported in Chapter 5 of this thesis. If listeners ignore the filtering effect of the plate, they will probably also not notice a change in this filtering.

Apparently some of the auditory information is interpreted while other data, though available, remains unidentified. Because the step of interpreting may or may not occur following the step of detection of the differences, we can say

that the perception of mechanical properties is a layered process. In the previous chapter we found a difference in steepness of the psychometric functions for the size and the speed judgments in a paired-comparison experiment, 2.6 and 1.7 respectively, as listed in Table 7.4. In this chapter we showed that the distances in the perceptual space for size and speed differences of the same sounds relate as 0.80 to 0.57, as listed in Table 8.2. We need to use the transformed perceptual space here, because we are interested in the audible differences caused by the speed and size separately, since they were judged separately in the tasks that resulted in the psychometric functions of Chapter 7. The ratios of the steepness, $2.6/1.7 = 1.5$ and the perceptual distance $0.80/0.57 = 1.4$. In the appendix of Chapter 7 we argued that twice as large perceptual distances lead to twice as large steepnesses of the psychometric function. Apparently, the distances in perceptual space underly the size and speed judgments of the listeners. Hence, judging differences in size and speed may be based on a low level of processing without the interference of cognitive processing. Stated differently, no evidence was found that listeners do “tune their ears” to listen for the size or speed of the rolling balls. The different performance of the estimation of the size and the speed of the rolling balls, as found in Chapter 7, can be explained by the detectability of differences in the sounds. One may then speculate whether the speed and size judgments are based on low-level perceptual attributes such as loudness, brightness, and roughness. Although these attributes will in combination play a significant role in size and speed judgment, it does not seem likely that only one of them will play a separate roll in, e.g., size and speed judgment. In other words, they are unlikely to function as independent auditory cues.

It is often suggested that in psychomechanics the properties of the mechanical system are mediated through the air to our brain which constructs an image of the mechanical system. This process is supposed to be unaffected by the path over which it is mediated, as, for instance, illustrated in the room acoustics that have no influence on the perception of the perceived mechanical properties. Our experiments indicate that the plate can be seen as part of the path from source to ear and that the properties of the path, here the plate, do not influence the perception of the mechanical properties of the ball. We can interpret this result as an example of the often suggested processes that listeners filter perceptually relevant mechanical characteristics of the acoustic event from other information in the sound signal. We are satisfied that the results we found illustrate the process that underlies this perceptive quality.

8.A Transformation of the MDS space

In the main text of this chapter we used a transformation to convert the three axes of an MDS space from directions without a specific meaning in terms of stimulus parameters into three meaningful ones, that is three directions that correlate most with the mechanical properties of the setup. In a vector, X_n we collect these mechanical differences in ball speed and size and in plate thickness, for each stimulus n :

$$X_n = \log \begin{pmatrix} \text{thickness} \\ \text{speed} \\ \text{size} \end{pmatrix}, \quad (\text{A-1})$$

where the logarithmic function is used to transform the differences in mechanical parameters into a ratio scale, because equally sized steps are perceived equally different, as was shown, at least for the speed and size part, in Chapter 7. For the thickness this has not been shown, but we will assume this works the same way. When we take just linear values, instead of the logarithmic ones, the conclusions presented in this appendix will not change, we will only change the numerical values, and if we want to compare these values with results in Chapter 7, we should take the logarithmic scales. The matrix that represents the mechanical properties X_n for each of the used setups will be denoted with X .

Using MDS we found the positions of the individual sounds, represented in the vector y , in the perceptual space. We will try to explain these positions by considering a linear combination of the differences in the mechanical system that produced the sound:

$$y = X\beta + \epsilon. \quad (\text{A-2})$$

We search $\hat{\beta}$ so that $\hat{y}$ best approximates y in a minimal sum of squares sense,

$$\hat{y} = X\hat{\beta}, \quad (\text{A-3})$$

from which $\hat{\beta}$ can be solved in the usual linear regression method,

$$\hat{\beta} = (X'X)^{-1}X'y. \quad (\text{A-4})$$

Next, we normalize this matrix,

$$\bar{\beta} = \frac{\hat{\beta}}{|\hat{\beta}|}, \quad (\text{A-5})$$

and take the inverse,

$$T = \bar{\beta}^{-1}. \quad (\text{A-6})$$

This transformational matrix T can be used to transform the positions found in the MDS space, y into positions y_{new} in a new space:

$$y_{\text{new}} = Ty. \quad (\text{A-7})$$

The axes correlating best with the mechanical parameter were not orthogonal in the original MDS space. They become orthogonal in the transformed MDS space. This means that the transformation is not orthogonal, and thus that the distances in the original space are not preserved in the transformed MDS space.

9 | Discussion

9.1 Introduction

In sound, many things are heard that can be qualified as mechanical properties of the objects producing the sound. We have tried to study the relation between the perceived object properties and the physical properties of the mechanical process creating the sound. We have not stopped after determining the perceived parameters alone, but we have also tried to get to know the processes underlying the perception of mechanical properties. Understanding these underlying processes has several advantages. When it is known how the properties are perceived:

- We are more certain that the perception is not based on artifacts or coincidentally available cues.
- We can try to predict perception results for situations not identical to but similar to the studied setup, e.g., for different materials or for different types of interaction, such as scraping instead of rolling.
- Efficient synthesis algorithms can be formulated, synthesizing only aspects that are used in the perception of the physical properties and hereby conveying information.
- We can determine the origin of an interaction in the perception of two mechanical features, for instance, size and speed of a rolling ball. The acoustic differences between the sounds generated by a fast and by a small ball could be inaudible to the listener, or, alternatively the listener could be unable to interpret the audible differences between the sounds.

In this chapter we will have a new look upon the schematic representation of our view on psychomechanics as presented in Figure 1.1 of the General Introduction presented in Chapter 1.

We have varied scalable properties of the mechanical system, and tested if these variations are perceivable. In the used setup these were: ball size, speed,

rolling direction and plate thickness. The resulting changes in the sound could or sometimes could not be detected by the listeners. Only variations in the sound that are detected can possibly be used to identify the variations of the mechanical system. The research can therefore be conducted in two steps. The first step is to determine the capability of listeners to *distinguish* two rolling sounds and the second step is to determine their performance in *identifying* the variations in the mechanical configuration.

9.1.1 Acoustics

We have contributed to the domain of acoustics in the form of providing a new numerical model for the generation of rolling ball sounds. We have also investigated the contact between ball and plate. This insight can be used not only for perception research, but also in different domains such as in the reduction of sound generation in roller bearings or by railroad trains, and to generate sound for virtual reality systems, when two virtual objects come in contact.

9.1.2 On interdisciplinary research

After all the many years of acoustics research, and despite important progress in this field, the most important questions from the point of view of psychomechanics are still unanswered. On the other hand the most important questions of psychoacoustic nature that are posed by physical acousticians are also unanswered. For instance, we do not know what makes a flute sound good when it sounds good. There is no computational algorithm that is able to determine from the sound if the played flute is a good or bad instrument. Yet for human listeners this is an easy task. Provided with this information, physical acousticians could turn a bad flute into a good one. Given the fact that both domains already exist for many decades, it is not likely that the psychoacousticians and the physical acousticians will answer each others questions soon. The only way to find the answer is to become skilled in both domains, in order to formulate the relevant questions and to find these answers. Therefore in this research both the mechanical modeling and the perceptual modeling of rolling and bouncing sounds have been treated in parallel.




FIGURE 9.1: Overview of the psychomechanical hearing process. For details see text. The indicated numbers show the chapters in the text that deal with the relation indicated with the arrow.

9.2 Psychomechanic research scheme

The relation between acoustics, mechanics and perception of events as they have been studied in this thesis is schematically presented in Figure 9.1. This figure gives an overview of the acoustical and perceptual processes that take place when a listener perceives the mechanical features of the objects producing a sound. It is not so much a view of the working of our ears and brains, but more a representation of the altering and reduction of information along this path. The content of the information tends to transform in each step along the line from the *mechanical system* at the upper left of Figure 9.1 towards the *perceived mechanical features* at the upper right of Figure 9.1. One of the aspects of this transformation is the change in dimensionality of this information as it passes from stage to stage. With dimensionality we mean the degrees of freedom of the parameter space describing each stage. Each independent parameter adds at least one dimension to this space; sometimes more than one, for instance, the size of the rectangular plate adds three dimensions, length, width and thickness.

We can compare this with the color space in the visual domain. We can make green light out of a combination of light rays that have a frequency corresponding to blue and yellow light, or using only light of one frequency corresponding to green light. If chosen right, we cannot see the difference. In the physical domain, light is described as a spectrum, having an amplitude for each frequency. In the perceptual domain color is uniquely identified by hue, brightness and saturation.

The upper left block of our scheme, labeled *mechanical system*, has a nearly infinite dimensionality, simply because we cannot tell where the mechanical system stops. The room in which an acoustic event takes place, for instance, has a profound influence on the sound produced by that event. In the upper right block, the dimensionality is rather low, since we mostly concentrate on one or two *perceived mechanical features*. We will now describe each stage separately.

Mechanical configuration, dynamics and material properties. There exists a vast amount of material properties. For instance, the elasticity of the ball and the plate, their sizes and densities, have a strong influence on the sound, as have the damping coefficients and the surface roughness parameters. To generate a sound via simulations we need to know all those parameters.

The resulting sound, however, can only vary in a limited number of ways. Indeed, if we take a simple example, the ideal string, we can change the pitch of the string by changing the string tension but also by changing its density. For an ideal string, if we multiply the string tension and density by an equal amount,

this leads to exactly the same oscillation frequency. In fact, also considering the string length, there are three parameters determining the resonance frequency. In less idealized systems we could try more complicated things; to compensate different viscoelastic damping properties of a string we can try, e.g., to change the driving force. For a physics based synthesis algorithm, such as our rolling ball algorithm, there are many parameters, but the acoustically important parameters, such as the frequencies of the vibrational modes of the plate are much lower in dimensionality than the mechanical parameters. In other words, changing one of the parameters of the sound generation process, either recorded or synthesized, does not change the sound in a unique way, because various mechanical variations result in comparable acoustic variations.

Mechanical basis variables. It is more efficient to consider only the resonance frequency of the string in stead of the density, length and tension of the string. For more complicated systems the *mechanical basis variables* ideally form a minimal set of variables that are orthogonal and form a basis for all the possible changes of the mechanical system that affect the generated sound. A very interesting technique to find these unique mechanical basis variables is dimensional analysis [13]. However, in the case of rolling sounds the dimensional analysis is very hard to do. The vast amount of parameters and the non-linear, non-continuous contact make the problem complex. Instead of using a mathematical technique to find the *mechanical basis variables*, we can resort to alternative variables that we derive by analyzing the mechanical system. A set of basis variables found this way may not be complete, orthogonal or the minimal possible set. But it can still be useful if a large reduction of the amount of parameters used to describe the material properties is found. For instance, in Chapter 2 we found two parameters for the bouncing ball mechanics, the inelasticity parameter λ and the contact time t_c , that govern the spectral and temporal variations of the generated sound.

Acoustic radiation. All properties of the sound pressure field in the air caused by the mechanical radiation fall under this category. Sound pressure level and frequency are examples of features in this domain. The effect of the acoustic radiation can be modeled by a linear filter with as input the vibrations of the object at the contact with the air and as output the air pressure at the ear of the listener. One interesting and ecologically very important phenomenon is that human listeners can detect and identify mechanical aspects of their environment almost independently of the acoustic environment. It is, therefore, likely that the features

in the acoustic signal that lead to our perception of the features of the mechanical system, are more or less independent of the acoustic transformation. We have, therefore, not concentrated our research on acoustic radiation.

Auditory information. All detectable variations *of a sound* can be considered as containing auditory information. In this sense, information is used as is usual in information theory. Hence, any variation of the sound that is detectable can possibly convey information. The perceptual attributes, pitch, loudness, and sharpness are examples of features in this domain. Much research has been devoted to the study of the processes underlying these perceptual attributes, see for instance Zwicker and Fastl [112] or Moore [70]. Computational models of these processes have been developed which can more or less accurately estimate the value the human listener gives to these perceptual attributes. The complexity of these models shows that there is no simple relation between perceptual attribute and an acoustical feature of the sound. They combine temporal information from a large number of auditory filters in order to come to a fair estimation of the perceived attribute. Nevertheless, these models up to now only operate accurately for a more or less limited set of relatively simple sounds. Certainly when the sound is generated by a combination of sound sources these models fail in attributing a plausible separate value to each of the perceived sound sources. These situations occur when, e.g., some musicians play together or some speakers talk simultaneously.

Moreover, we do not yet know whether these perceptual attributes fully characterize the sound. According to the ANSI standard [1] there are only four perceptual attributes: pitch, loudness duration and timbre. Hence, well established attributes such as roughness and sharpness are timbral attributes. There probably are quite a lot of timbral attributes of a sound that on the one hand are used for the perception of some mechanical parameter but to which no well established name has yet been given. Moreover, when people are asked to describe a mechanical sound they not only use timbral attributes but also refer to the perceived physical properties of the sound source such as the material of the sound producing object and the kind of contact they perceive. In addition, they may use onomatopoeic descriptions. As a consequence, it is difficult to get to know the structure of the auditory information present in mechanical sounds on the basis of verbal descriptions or magnitude estimations. So in this thesis the problem was addressed from another side.

Experimental methodologies such as difference scaling, where people are not asked for any meaning or perceptual attribute of the sound, but are only tested in their capability to differentiate sounds, were used to investigate the structure of the *auditory information*. We described this method in Chapter 8, and by using multidimensional scaling on these data, we estimated the number of dimensions of the space of auditory information for the rolling sounds used throughout this thesis.

Auditory cues. All detectable variations *in the mechanical system* may provide us with cues used in perceiving the mechanical features. In contrast to the *auditory information*, the *auditory cues* give meaning to the differences that are heard. By signal processing techniques we can easily make detectable changes to the sound that do not have any meaning. Sound that is manipulated by distorting it or by adding easily detected noise does not provide us with information about the mechanical system. Thus, by this manipulation the *auditory information* is changed but not necessarily the *auditory cues*. The inverse is not possible. If two sounds are associated with, for instance, two rolling balls with different sizes, there must be a detectable difference in the sound and thus they must differ in *auditory information*.

Previously, the cues were considered apriori and by the imagination of the researcher. Then, by a correlation of the perceived properties of the objects, i.e., size and speed in our case, with the presumed auditory cues, such as the sharpness of the sound, it was tried to determine to which extent these auditory cues were used by the listener to determine the object properties. This approach was successful in determining some of the auditory cues but lacks precision in positioning these cues. If we investigated cue A, which correlates with some perceived *mechanical feature*, we still do not know if there exists another cue B that might correlate better with this perceived *mechanical feature*. Thus, previously sharpness and amplitude modulation were considered as cues [40] but probably more cues can be used. By examining this space we can determine how much auditory information the listener uses in extracting the auditory cues used in detecting the mechanical features.

Note that the classical perceptual qualities such as pitch, loudness and sharpness surely belong to the *auditory information* or can be represented by a combination of other *auditory information* variables, since perception experiments have shown that people can detect changes in these percepts. They could also be part of the *auditory cues* if they convey information about the mechanical system to

the listener. Other *auditory cues* might be the result of a combination of different parameters of the auditory information. Some *auditory information* might not be used at all as *auditory cue*, even if it provides meaningful information about the mechanical system. An example of this we saw in Chapter 8 where we found one of the dimensions of the space resulting from the MDS experiment that correlated directly with the plate thickness. But the plate thickness was not among the perceived mechanical features as we saw in Chapter 7. Another example is the well known effect, the stability of loudness [78] also known as loudness constancy [110], which states that the perceived loudness of, for instance, a vocal sound, is not based on the signal power of the sound at our ears. In fact, if the sound pressure of a sound gets less we will in general perceive it as more distant and not as less loud.

It is questionable whether auditory features like roughness or sharpness are the only auditory cues used in the perception of one specific physical parameter. Indeed, let this be true and let, e.g., the sharpness be a cue for the size of the ball where sharper sounds are related to smaller balls. Then changing the sharpness in any way, and leaving the other auditory cues unaltered would lead to the percept of a smaller ball. It is questionable whether this sound is a rolling sound, i.e., if there can exist a system that would radiate such a sound. Stated differently, for a given rolling sound there might be another rolling sound that only differs in, say, sharpness. If there is such a rolling sound, it is probable that all or nearly all mechanical parameters need to be changed to generate this sound.

In the experiments described in Chapter 7 where participants were asked to judge the absolute size of rolling balls after listening to their sounds, no structural interaction was found between the speed and size judgments. Thus if sharpness is a cue for both size and speed perception, there must be a second variable to decorrelate the two.

Perceived mechanical features are mechanical properties of the system as they are perceived by the listener, such as the size and the speed of the ball. All detectable and identifiable variations *in the mechanical system* are represented in the *auditory cues*, by definition. The listener uses them to build a mental image of the sound generating system, that is, its mechanical properties such as the size of the ball and the plate, and the type and speed of the interaction between the ball and plate.

9.2.1 Experiments

Nearly all our perception experiments were carried out with recorded sounds, the exception being described in Chapter 5. Only in Chapter 3 we manipulated the recorded sounds. In Figure 9.1, depicting the various stages in perceiving an acoustic event, we indicate for each experiment the chapter in which this experiment is described and which relation within the scheme is investigated. In the experiment using the absolute scaling paradigm described in Chapter 7, we determined the accuracy of the size and speed judgments of people listening to rolling sounds. In terms of the scheme, we determined the relation between the *mechanical system* and the *perceived mechanical features* with this experiment. In the MDS experiment with recorded sounds, as discussed in Chapter 8, we determined the relation between the *mechanical system* and the *auditory information*. By comparing the results of these two experiments, we have also derived some relations between the *auditory information* and the *perceived mechanical information*. In Chapter 2 we showed the relation of two *mechanical basis variables*, the inelasticity parameter λ and the contact time t_c , and the spectrum and the temporal structure of the bouncing sound. This analysis helped to understand that both spectrum and temporal structure can be used to determine the size of the bouncing ball. We showed in Chapter 3 that people rely on either the spectrum or the temporal structure of the bouncing sound, or on both to judge the size of a bouncing ball, a feature represented in the *perceived mechanical feature* block of our scheme. In Chapter 5 we have shown that acoustic information that at first sight may seem to be a good candidate for being an auditory cue, the spectro-temporal effect caused by the reflections of the vibrational waves at the edges of the plate on the rolling sound, was in fact not used by listeners.

9.2.2 Inferred relations between blocks

In the left hand side of the Figure 9.1, representing physical reality, we can transform variables from one block of the mechanical system to another block by using mechanical formulas. In signal processing terms we call this the transfer function. We have studied the acoustics of bouncing and of rolling sounds in Chapters 2 and in Chapters 4 and 6, respectively. To find the transfer function from one block on the left to one block on the right we must conduct perception experiments. However, we cannot directly find the relation between blocks at the right hand side. All *auditory information* that is used by the listener to extract *perceived mechanical features* must be represented in the *auditory cues*. Hence, the relation

between blocks on the right can be found by analyzing the results obtained in two perception experiments, as we did by combining the results of experiments described in Chapters 7 and 8.

9.2.3 Synthesized sounds

The main advantage of using synthesized sounds over using recorded sounds is that we can more easily explore all possible sounds within a certain physical parameter space. For recorded sounds, a small slow ball resembles a large fast ball, and some combination on a thick plate can resemble another on a thin plate. Finding the more different sounds, in physical terms, is easier if we can modify the *mechanical basis variables* directly, which is normally only possible if we use synthesized sound. Moreover, it is possible to vary the synthesized sounds using only a few values for all the basis variables, and all sounds will physically be quite different. Because there are relatively few *mechanical basis variables* compared to possible variations in the *mechanical configuration, dynamics and material properties*, it is much more practical to use the *mechanical basis variables* to explore the space of the *auditory information* and that of *auditory cues*. Using recorded sounds, and varying the size, speed and plate thickness stepwise, we can obtain many sounds that are physically more or less similar. Their difference ratings in the perception experiments can be expected to be quite low, and this will provide us with little information about the auditory processing of these sounds.

Another unique possibility of a synthesis system is to generate new, unnatural, situations where the auditory cues give contradictory information. By again asking the participants for the perceived size and velocity of the ball we could determine which cues are perceptually given more weight than others and, hence, measure the salience of the various cues. We have applied this method to the sound of bouncing balls, where we could use manipulated recordings of bouncing balls instead of synthetic sounds, because it was possible to manipulate bouncing sounds to have independent temporal and spectral properties for the size of the ball. Such an approach was feasible for the bouncing sounds because of the silent gaps between the bounces of the ball. For rolling sounds, that are continuous, we will need to generate the sounds synthetically to use this approach. We believe that the use of sounds that are synthesized in this way, helps to overcome some of the problems that exist in similar research using manipulated recorded sounds. Such research using manipulated recorded sounds was done by Houben [40], and the difficulties were that one temporal effect, amplitude modulation, could not be removed from the sounds and to vary this tempo-

ral effect, it was added in a perhaps exaggerated way with different periodicities. In another experiment he combined the temporal pattern from one rolling sound with the spectral properties of another rolling sound. This manipulation leads, however, to audible artifacts in the sound. Furthermore, as long as we don not know what the temporal cues for the speed and size of a rolling ball exactly are, we also don not know how effective the procedure is in manipulating these temporal cues. In a synthesis algorithm, we know what parameters influence the temporal structure of the sound, and we can establish the relation between these parameters and the perceived size and speed.

An advantage of using the mechanical basis variables to synthesize unrealistic sounds over changing sounds with signal processing is that the resulting sounds are more natural in the sense that they obey physical laws. When we change a sound using signal processing, we obtain another sound that most probably could not be produced by another ball-plate configuration. For instance, if we change the spectral properties of a rolling sound to have more high frequencies we obtain a sound that has probably no mechanical representation. We can test how this alters the *perceived mechanical parameters* but we cannot tell if this reflects the new *mechanical configuration* in a correct way. Once we have identified the *mechanical basis variables* that change the high-frequency content of a sound, we can alter those in our synthesis algorithm. By using these sounds in a perception test we can again test how this alters the *perceived mechanical parameters*, but this time we can track back what changes in the *mechanical configuration* were required to create these sound changes. This *mechanical configuration* may or may not exist, but if the values of the *mechanical basis variables* do not deviate too much from real values, the sound should not be perceived as unnatural. As an example, if we drink a beverage with the color of carrot juice and the taste of coffee this would be perceived as unnatural. But if we slightly change the processing of the beans this would not likely change the recognition of the coffee.

Thus, to summarize the last few paragraphs, different acoustic properties of a sound generating system are very often correlated. For instance the amplitude-modulation frequency depends on the size and the speed of the rolling ball, as does, in a different way, the average spectrum of the rolling sound. We argued that, by using a synthesis algorithm based on the mechanical basis variables, we can in principle still fulfill this relation, and generate new rolling sounds that are more natural than is the case for manipulated recording sounds. By using the mechanical basis variables it should be easier to find mechanical setups with which one may generate sounds that differ as much as possible from an acoustical point

of view. These setups may be difficult to realize in practice due to, e.g., the impossibility to produce some materials and shapes with a specified combination of mechanical properties. In these cases we can use the synthesis algorithm to generate these sounds. Alternatively, we can synthesize very different sounds such that the quantities that are normally strongly correlated, are varied independently. To do so, we may analyze the mechanical basis variables and split one of them into two variables. In this way we create a degree of freedom that was not present in the real mechanical setup. For instance, let us consider the radius of the ball. It has at least three influences in the synthesis of rolling sounds. First, in the Hertz contact law, second, in the periodicity of the driving force of the ball due to imperfect sphericity of the ball, and third, in the weight of the ball. These are thus all governed by one mechanical basis variable. By splitting it in two we can generate sounds that have the quasi-periodicity of a large ball but the average spectrum of a small ball. We actually used such an approach in the study of the perception of bouncing balls in Chapter 3, because for bouncing balls, separation of temporal and spectral properties could be achieved without a synthesis algorithm. Houben [40] used such an approach by manipulating the spectral and temporal properties of recorded rolling sounds, but here the resynthesized sounds were perceptually different from the recorded sounds. The results obtained by Houben were valid since he controlled for this difference. But, in contrast to Houben's approach, by synthesizing sounds based on the mechanical basis variables we obtain sounds with properties that can be tracked down to their mechanical characteristics. By then presenting listeners with these sounds, and asking them for the perceived mechanical features, we can in principle track these perceived features down to the real mechanical features with as valid a stimulus set as is acoustically possible. We have thus shown the viability for such research methods, and have taken the first steps with our synthesis algorithm for rolling sounds. Naturally more research is needed to exploit the full potential of such a paradigm.

Bibliography

- [1] American national psychoacoustical terminology. s3.20. American Standards Association, 1973.
- [2] M. Abramowitz and I. A. Stegun, editors. *Handbook of Mathematical Functions, with Formulas, Graphs, and Mathematical Tables*. Courier Dover Publications, 1965.
- [3] P. Allen and C.-A. Bond. Multidimensional scaling of complex sounds by school-aged children and adults. *J. Acoust. Soc. Am.*, 102(4):2255–2263, 1997.
- [4] G. B. Arfken and H. J. Weber. *Mathematical Methods for Physicists*. Academic Press, 2000.
- [5] J. S. Bendat and A. G. Piersol. *Engineering Applications of Correlation and Spectral Analyses*. Wiley, 1993.
- [6] F. A. Bilsen. Repetition pitch: Monaural interaction of a sound with the repetition of the same, but phase shifted, sound. *Acustica*, 17:295–300, 1966.
- [7] A. S. Bregman. *Auditory Scene Analysis: The Perceptual Organization of Sound*. The MIT Press, 1990.
- [8] N. Brilliantov and T. Pöschel. Rolling as a “continuing collision”. *The European Physical Journal B - Condensed Matter and Complex Systems*, 12(2):299–301, 1999.
- [9] P. A. Cabe and J. B. Pittenger. Human sensitivity to acoustic information from vessel filling. *Journal of Experimental Psychology: Human Perception and Performance*, 26:313–324, 2000.
- [10] G. Canévet, D. Habault, S. Meunier, and F. Demirdjian. Auditory perception of sounds radiated by a fluent-loaded vibrating plate excited by a point force. *Acta Acustica united with Acustica*, 90:181–193, 2004.

- [11] A. Chaigne and V. Doutaut. Numerical simulations of xylophones. I. Time-domain modeling of the vibrating bars. *J. Acoust. Soc. Am.*, 101(1):539–557, 1997.
- [12] A. Chaigne and C. Lambourg. Time-domain simulation of damped impacted plates. I. Theory and experiments. *J. Acoust. Soc. Am.*, 109(4):1422–1432, 2001.
- [13] A. Chaigne and C. Stoelinga. Simulation numérique de sons de roulements. In *Proceedings of the 8ème Congrès Français d’Acoustique, Tours 24-27 Avril 2006*, 2006.
- [14] L. Cremer, M. Heckl, and E. E. Ungar. *Structure-Borne Sound*. Springer-Verlag, second edition, 1988.
- [15] R. Cross. The bounce of a ball. *Am. J. Phys.*, 67(3):222–227, 1999.
- [16] F. Demirdjian, D. Habault, S. Meunier, and G. Canévet. Can we hear the complexity of vibrating plates? In *Proceedings of the CFADAGA conference, Strasbourg*, 2004.
- [17] P. Devillard. Dynamics of a ball on a vibrating table. *J. Phys. I France*, 4:1003–1011, 1994.
- [18] A. P. Dowling and J. E. F. Williams. *Sound and Sources of Sound*. Ellis Horwood Limited, 1983.
- [19] E. Eichler. Plate-edge admittances. *J. Acoust. Soc. Am.*, 36:344–348, 1964.
- [20] E. Falcon, C. Laroche, S. Fauve, and C. Coste. Behavior of one inelastic ball bouncing repeatedly off the ground. *Eur. Phys. J. B*, 3:45–57, 1998.
- [21] J. Feldmann. Point distributed static load of rough elastic contact can cause broadband vibrations during rolling. *Mechanical Systems and Signal Processing*, 16(2–3):285–302, 2002.
- [22] D. J. Freed. Auditory correlates of perceived mallet hardness for a set of recorded percussive sound events. *J. Acoust. Soc. Am.*, 87(1):311–322, 1990.
- [23] W. W. Gaver. How in the world do we hear? Explorations of ecological acoustics. *Ecological Psychology*, 5:285–313, 1993.
- [24] W. W. Gaver. What in the world do we hear? An ecological approach to auditory event perception. *Ecological Psychology*, 5:1–29, 1993.

- [25] G. A. Gescheider. *Psychophysics: The Fundamentals*. Lawrence Erlbaum Associates Inc, 1997.
- [26] J. J. Gibson. *The senses considered as perceptual systems*. Houghton Mifflin, 1966.
- [27] J. M. Gilchrist, D. Jerwood, and H. Samisrael. Comparing and unifying slope estimates across psychometric function models. *Perception & Psychophysics*, 67(7):1289–1303, 2005.
- [28] B. L. Giordano. *Sound Source Perception in Impact Sounds*. PhD thesis, università degli studi di padova, 2005.
- [29] W. Goldsmith. *Impact: The Theory and Physical Behaviour of Colliding Solids*. Dover Publications, 2001.
- [30] K. F. Graff. *Wave Motion in Elastic Solids*. Dover Publications, New York, 1991.
- [31] M. Grassi. Do we hear size or sound? Balls dropped on plates. *Perception & Psychophysics*, 67(2):274–284, 2005.
- [32] J. A. Greenwood and J. B. P. Williamson. Contact of nominally flat surfaces. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 295(1442):300–319, 1966.
- [33] J. M. Grey. Multidimensional perceptual scaling of musical timbres. *J. Acoust. Soc. Am.*, 61(5):1270–1277, 1977.
- [34] J. Guckenheimer and P. Holmes. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer-Verlag, New York, 1983.
- [35] J. H. D. Guimarães. Equivalent surface approach in modeling of rough elastic contact in rolling bearings. In *Proceedings of the DAGA*, pages 338–339, 2003.
- [36] D. J. Hermes. Synthesis of the sounds produced by rolling balls. Technical report, IPO Rapport no. 1226, 2000.
- [37] H. Hertz. Ueber die Berührung fester elastischer Körper. *Zeitschrift für die Reine und Angewandte Mathematik*, pages 156–171, 1881.
- [38] D. P. Hess and A. Soom. Normal vibrations and friction under harmonic loads: Part I – Hertzian contacts. *Transactions of the ASME*, 113:80–86, 1991.

- [39] M. Houben, D. Hermes, and A. Kohnrausch. Auditory perception of the size and velocity of rolling balls. *IPO Annual Progress Report*, 34:86–93, 1999.
- [40] M. M. J. Houben. *The Sound of Rolling Objects: Perception of Size and Speed*. PhD thesis, J.F. Schouten School for User-System Interaction Research, Eindhoven University of Technology, 2002.
- [41] M. M. J. Houben, A. Kohlrausch, and D. J. Hermes. Auditory cues determining the perception of the size and speed of rolling balls. In *Proceedings of the 2001 International Conference on Auditory Display*, Helsinki, Finland, 2001.
- [42] M. M. J. Houben, A. Kohlrausch, and D. J. Hermes. Perception of the size and speed of rolling balls by sound. *Speech communication*, 43:331–345, 2004.
- [43] M. M. J. Houben, A. Kohlrausch, and D. J. Hermes. The contribution of spectral and temporal information to the auditory perception of the size and speed of rolling balls. *Acta Acustica united with Acustica*, 91:1007–1015, 2005.
- [44] M. M. J. Houben and C. N. J. Stoelinga. Some temporal aspects of rolling sounds. In *Journées Design Sonore 2002*, Paris, France, 2002.
- [45] S. Hyun, L. Pei, J. F. Molinari, and M. O. Robbins. Finite-element analysis of contact between elastic self-affine surfaces. *Phys. Rev. E*, 70(2):026117, 2004.
- [46] K. L. Johnson. *Contact Mechanics*. Cambridge University Press, 1987.
- [47] P. Joly and L. Rhaouti. A qualitative analysis of a simplified model for the non linear membrane-mallet interaction. Technical Report 3234, INRIA, 1997.
- [48] M. C. Junger and D. Feit. *Sound, Structures, and Their Interaction*. MIT Press, 1984.
- [49] T. Kaczmarek. Auditory perception of sound source velocity. *J. Acoust. Soc. Am.*, 117(5):3149–3156, 2005.
- [50] B. C. Kirkwood. The influence of presentation method on auditory length perception. *Studies in Perception and Action*, pages 145–148. Lawrence Erlbaum Associates, 2005.

- [51] L. Kogut and I. Etsion. An improved elastic-plastic model for the contact of rough surfaces. In *3rd Aimeta International Tribology Conference*, Italy, 2002.
- [52] L. Kogut and L. Etsion. Elastic-plastic contact analysis of a sphere and a rigid flat. *Journal of Applied Mechanics*, 69:657–662, 2002.
- [53] E. Kreyzig. *Advanced Engineering Mathematics*. Wiley, eighth edition, 1999.
- [54] A. J. Kunkler-Peck and M. Turvey. Hearing shape. *Journal of Experimental Psychology: Human Perception and Performance*, 26(1):279–294, 2000.
- [55] G. Kuwabara and K. Kono. Restitution coefficient in a collision between two spheres. *Jpn. J. Appl. Phys.*, 26(8):1230–1233, 1987.
- [56] S. Lakatos, S. McAdams, and R. Caussé. The representation of auditory source characteristics: Simple geometric form. *Perception & Psychophysics*, 59(8):1180–1190, 1997.
- [57] C. Lambourg, A. Chaigne, and D. Matignon. Time-domain simulation of damped impacted plates, II. Numerical model and results. *J. Acoust. Soc. Am.*, 109(4):1422–1432, 2001.
- [58] L. D. Landau and E. M. Lifshitz. *Theory of Elasticity*. Butterworth-Heinemann, 1986.
- [59] D. N. Lee and P. E. Reddish. Plummeting gannets: a paradigm of ecological optics. *Nature*, 293:293–294, 1981.
- [60] X. Li, R. J. Logan, and R. E. Pastore. Perception of acoustic source characteristics: Walking sounds. *J. Acoust. Soc. Am.*, 90(6):3036–3049, 1991.
- [61] J. M. Luck and A. Mehta. Bouncing ball with a finite restitution: Chattering, locking, and chaos. *Phys. Rev. E*, 48(5):3988–3997, 1993.
- [62] R. A. Lutfi. Informational processing of complex sound: I. Intensity discrimination. *J. Acoust. Soc. Am.*, 86(3):934–944, 1989.
- [63] R. A. Lutfi. Auditory detection of hollowness. *J. Acoust. Soc. Am.*, 110(2):1010–1019, 2001.
- [64] R. A. Lutfi, K. A. Doherty, and E. Oh. Psychometric functions for the discrimination of spectral variance. *J. Acoust. Soc. Am.*, 100(4):2258–2265, 1996.

- [65] J.-B. Martens. Multidimensional modeling of image quality. *Proceedings of the IEEE*, 90(1):133–153, 2002.
- [66] S. McAdams. The psychomechanics of real and simulated sound sources. *J. Acoust. Soc. Am.*, 107(5):2792, 2000.
- [67] S. McAdams and E. Bigand, editors. *Thinking in Sound, The cognitive Psychology of Human Audition*. Clarendon Press, 1993.
- [68] S. McAdams, A. Chaigne, and V. Roussarie. The psychomechanics of simulated sound sources: Material properties of impacted bars. *J. Acoust. Soc. Am.*, 115(3):1306–1320, 2004.
- [69] R. D. Mindlin. influence of rotary inertia and shear on flexural motions of isotropic, elastic plates. *Journal of Applied Mechanics*, pages 31–38, 1951.
- [70] B. C. J. Moore. *An Introduction to the Psychology of Hearing*. Academic Press, 5th edition, 2003.
- [71] P. M. Morse and K. U. Ingard. *Theoretical Acoustics*. McGraw–Hill, 1968.
- [72] P. R. Nayak. Contact vibrations. *Journal of Sound and Vibration*, 22(3):297–322, 1972.
- [73] A. H. Nayfeh and D. T. Mook. *Nonlinear Oscillations*. John Wiley & Sons, Inc, Wiley classics library edition, 1995.
- [74] A. Nordborg. Wheel/rail noise generation due to nonlinear effects and parametric excitation. *J. Acoust. Soc. Am.*, 111(4):1772–1781, 2002.
- [75] M. Pärssinen. Hertzian contact vibrations under random external excitation and surface roughness. *Journal of Sound and Vibration*, 214(4):779–783, 1998.
- [76] M. Pauly, D. K. Pai, and L. J. Guibas. Quasi-rigid objects in contact. In R. Boulic and D. K. Pai, editors, *Proceedings of the Eurographics/ACM SIG-GRRAPH Symposium on Computer Animation*, 2004.
- [77] S. Philips, L. Atlas, and J. Pitton. Automatic feature identification using multidimensional scaling data (abstract). *J. Acoust. Soc. Am.*, 119(5):3395, 2006.
- [78] R. Plomp. *The intelligent ear*. Lawrence Erlbaum Associates, 2002.
- [79] B. Porat. *A Course in Digital Signal Processing*. Wiley, 1997.

- [80] M. Rath and D. Rocchesso. Continuous sonic feedback from a rolling ball. *IEEE Multimedia*, 12(2):60–69, 2005.
- [81] M. Rath, D. Rocchesso, and F. Avanzini. Physically based real-time modeling of contact sounds. In *Proc. Int. Computer Music Conf.*, 2002.
- [82] T. W. Reese. The application of the theory of physical measurement to the measurement of psychological magnitudes, with three experimental examples. *Psychological Monographs*, 55:1–89, 1943.
- [83] E. Rigaud and J. Perret-Liaudet. Experiments and numerical results on non-linear vibrations of an impacting Hertzian contact. Part 1: Harmonic excitation. *Journal of Sound and Vibration*, 265(2):289–307, 2003.
- [84] D. Rocchesso and F. Fontana, editors. *The sounding object*. Edizioni di Mondo Estremo, 2003.
- [85] B. Russell. *The Principles of Mathematics*. Norton, New York, 1938.
- [86] J. Sabot, P. Krempf, and C. Janolin. Non-linear vibrations of a sphere-plane contact excited by a normal load. *Journal of Sound and Vibration*, 214(2):359–375, 1998.
- [87] R. S. Sayles and S. Y. Poon. Surface topology and rolling element vibration. *Precision Engineering*, 3(3):137–144, 1981.
- [88] J. Scheuren. Noise control for cars and trains from an engineering perspective. In *InterNoise conference proceedings*, Nice, France, 2000.
- [89] M. R. Schroeder. Synthesis of low-peak-factor signals and binary sequences with low autocorrelation. *IEEE Transactions on Information Theory*, 16:85–89, 1970.
- [90] R. Sondergaard, K. Chaney, and C. E. Brennen. Measurements of solid spheres bouncing off flat plates. *Journal of Applied Mechanics*, 57:694–699, 1990.
- [91] I. Stensgaard and E. Lægsfaard. Listening to the coefficient of restitution - revisited. *Am. J. Phys.*, 69(3):301–305, 2001.
- [92] S. S. Stevens. *Mathematics, Measurement and Psychophysics*, chapter 1, pages 1–49. John Wiley and sons, New York, 1951.

- [93] C. Stoelinga and A. Chaigne. Simulations and experiments regarding rolling objects. In *Proceedings of the Forum Acusticum conference*, pages 63–68, 2005.
- [94] C. Stoelinga, D. Hermes, and A. Kohlrausch. Using multidimensional scaling to determine distances in the perceptual space of rolling sounds (abstract). *J. Acoust. Soc. Am.*, 119(5):3237, 2006.
- [95] C. Stoelinga, D. Hermes, and A. Kohlrausch. Using psychometric functions to determine auditory perception of the size and speed of a rolling ball (abstract). *J. Acoust. Soc. Am.*, 119(5):3333, 2006.
- [96] C. Stoelinga, H. Peeters, and I. Lopez. The bouncing of balls: Influence of temporal and spectral changes on perception. In *Proceedings of the CFADAGA conference 22-25 march 2004, Strasbourg*, 2004.
- [97] C. N. J. Stoelinga and A. Chaigne. Time-domain modeling and simulation of rolling objects. Accepted for publication in *Acta Acustica united with Acustica*.
- [98] C. N. J. Stoelinga, D. J. Hermes, A. Hirschberg, and A. J. M. Houtsma. Temporal aspects of rolling sounds: A smooth ball approaching the edge of a plate. *Acta Acustica united with Acustica*, 89:809–817, 2003.
- [99] P. Swistak. Paradigms of measurement. *Theory and desicion*, 29:1–17, 1990.
- [100] D. J. Thompson, P. Fodiman, and H. Maé. Experimental validation of the twins prediction program for rolling noise, Part 2: Results. *Journal of Sound and Vibration*, 193(1):137–147, 1996.
- [101] D. J. Thompson, B. Hemsworth, and N. Vincent. Experimental validation of the twins prediction program for rolling noise, Part 1: Description of the model and method. *Journal of Sound and Vibration*, 193(1):123–135, 1996.
- [102] L. L. Thurstone. A law of comparative judgement. *Psychological Review*, 34:278–286, 1927.
- [103] W. S. Torgerson. *Theory and Methods of Scaling*. Wiley, 1958.
- [104] J. A. Turner. Non-linear vibrations of a beam with cantilever-hertzian contact boundary conditions. *Journal of Sound and Vibration*, 275(1-2):177–191, 2004.

- [105] K. van den Doel, P. G. Kry, and D. K. Pai. Foleyautomatic: Physically-based sound effects for interactive simulation and animation. In *SIGGRAPH conference proceedings*, 2001.
- [106] R. M. Warren. *Auditory Perception: A New Analysis and Synthesis*. Cambridge university press, 1999.
- [107] T. X. Wu and D. J. Thompson. Theoretical investigations of wheel/rail non-linear interaction due to roughness excitation. Technical Memorandum 852, Institute of Sound & Vibration Research, University of Southampton, 2000.
- [108] W. A. Yost. Strength of the pitches associated with ripple noise. *Journal of the Acoustical Society of America*, 64(2):485–492, 1978.
- [109] W. A. Yost. The dominance region and ripple noise pitch: A test of the peripheral weighting model. *Journal of the Acoustical Society of America*, 72(2):416–425, 1982.
- [110] P. Zahorik and F. L. Wightman. Loudness constancy with varying sound source distance. *Nature neuroscience*, 4(1):78–83, 2001.
- [111] C. Zener. The intrinsic inelasticity of large plates. *Phys. Rev.*, 59:669–673, 1941.
- [112] E. Zwicker and H. Fastl. *Psychoacoustics*. Springer-Verlag, 1990.

Summary

A Psychomechanical Study of Rolling Sounds

In this thesis both the mechanics and the perception of the sounds from rolling and bouncing objects are studied. It is necessary to combine these two research domains, resulting in what is called psychomechanics, because only in this combination can we discover the acoustical origin of the information a listener uses to detect mechanical properties of objects when listening to the sounds generated by these objects. In this thesis objects are balls rolling over or bouncing on a plate. The characteristics of such interactions are defined by physical properties of the components in the interaction, among them the size and the speed of the rolling balls, and, perhaps, the thickness of the plate over which the ball rolls. The research question is to what extent the listener can "hear" these physical properties and on what kind of information in the sound this is based. The first two chapters of this thesis investigate bouncing while the remaining five chapters deal with rolling sounds.

A typical parameter in investigating bouncing processes is the restitution coefficient, defined as the proportion of two subsequent time intervals between bounces. The traditional model for calculating the restitution coefficient of a ball was adopted by calculating the plate response from the point impedance of the plate. This analytical model was compared with a numerical model and with experimental results. The restitution coefficient is one source of information the listener may use to detect the size of the bouncing ball; the spectral content of each impact sound is another. The influence and relative weight of each of these two sources of information was investigated in a perception experiment.

The numerical model for simulating a single impact of a ball on a plate was adapted to simulate the movement and the sound of a ball rolling over a plate. The main changes were made in the contact between the ball and the plate. Instead of an impact point that was fixed in space and short in time, the model now incorporates a contact point that is variable in space and continuous in time.

Furthermore, the model was extended to contain the surface roughness of the plate. The validity of the model was demonstrated by simulating the ball and plate movement, and comparing the results with experiments and analytical calculations.

In the case of a ball moving over a plate, a temporal effect was observed due to the interference between the sound directly generated at the point of contact between ball and plate, and the sound reflected at the edge of the plate. This effect was added to synthesized rolling sounds which resulted in a more natural sound. In a related perception experiment it was shown that the acoustic differences in the sounds of balls rolling towards or away from the edge of a plate are sufficient to perceptually discriminate these sounds. Listeners were, however, not able to indicate the rolling direction of the ball. That is, listeners could not interpret the acoustic information correctly in terms of perceived mechanical features.

The auditory capabilities to perceive the size and speed of a rolling ball were studied in a difference-scaling task, where the participants were asked to judge the difference between two sounds without having to label the origin of these differences. In additional experiments participants were asked to rate explicitly the mechanical properties of ball and plate. The experimental methodologies used were paired comparison and absolute magnitude estimation. From these results the distances in the perceptual space of rolling sounds were derived. Furthermore, by calculating the psychometric function for the participants' estimations, an exact measure for the probability that a listener correctly identifies the difference in one of the physical parameters was calculated.

The influence of some object properties, here notably the plate thickness, can be detected in a difference scaling task, but listeners can not estimate the thickness in a magnitude scaling task. This indicates that the auditory perception of the mechanical object properties is a layered process, and that detection and estimation find place at a different level of such a layered process. On the other hand, for the two remaining mechanical properties, the size and the speed of the ball, the estimation of the size was more robust, in the sense that it was little influenced by a change of experimental methodology or simultaneous variation of the other mechanical parameters.

Curriculum Vitae

- 1974 Born in Eindhoven, the Netherlands
- 1986-1992 HAVO, Rommert Casimir, Eindhoven
- 1992-1997 Technische Hogeschool Eindhoven, B.Sc. in Electrical Engineering
- 1997-2001 Technische Universiteit Eindhoven, M.Sc. in Electrical Engineering
- 2001-2002 Technische Universiteit Eindhoven, Research assistant
- 2002-2007 Technische Universiteit Eindhoven, Ph.D. Student at the J.F.Schouten
School for User-System Interaction Research